

(12)
**United States Patent**  
**Ye et al.**

(10) **Patent No.:**      **US 9,171,073 B1**  
(45) **Date of Patent:**      **Oct. 27, 2015**

(54)	<b>INDEX MECHANISM FOR REPORT GENERATION</b>	7,548,928 B1	6/2009	Dean et al.	
		7,865,495 B1 *	1/2011	Roizen et al.	707/708
		2003/0212694 A1	11/2003	Potapov et al.	
		2004/0034616 A1	2/2004	Witkowski et al.	
		2005/0120006 A1 *	6/2005	Nye	707/3
		2007/0050384 A1 *	3/2007	Whang et al.	707/100
		2008/0040348 A1	2/2008	Lawande et al.	
		2008/0294676 A1	11/2008	Faerber et al.	
		2009/0144235 A1	6/2009	Bhide et al.	
		2009/0313210 A1	12/2009	Bestgen et al.	
		2010/0030796 A1	2/2010	Netz et al.	
		2010/0088309 A1	4/2010	Petculescu et al.	
		2010/0088315 A1	4/2010	Netz et al.	
		2010/0088352 A1	4/2010	Hendrickson et al.	

(21) Appl. No.: **13/618,192**

(22) Filed:      **Sep. 14, 2012**

**Related U.S. Application Data**

(63) Continuation of application No. 12/777,643, filed on May 11, 2010, now Pat. No. 8,452,755.

(60) Provisional application No. 61/177,489, filed on May 12, 2009.

(51) **Int. Cl.**  
**G06F 17/30**                      (2006.01)

(52) **U.S. Cl.**  
CPC .... **G06F 17/30722** (2013.01); **G06F 17/30011** (2013.01)

(58) **Field of Classification Search**  
CPC ..... G06F 17/30722; G06F 17/30011  
USPC ..... 707/706–709, 711, 716–719, 721, 723, 707/736, 758, 781, 999.001–999.003  
See application file for complete search history.

(56)                      **References Cited**

U.S. PATENT DOCUMENTS

5,603,022 A	2/1997	Ng et al.
5,678,043 A	10/1997	Ng et al.
5,918,225 A	6/1999	White et al.
6,594,678 B1	7/2003	Stoutamire et al.

**OTHER PUBLICATIONS**

United States Patent and Trademark Office; Non-final Office Action for U.S. Appl. No. 12/777,631 dated Mar. 22, 2012, 22 pages.  
U.S. Non-Final Office Action for U.S. Appl. No. 12/777,643 dated Jan. 3, 2012, 28 pages.

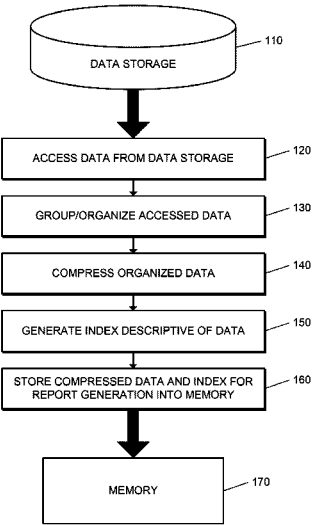
(Continued)

*Primary Examiner* — Syling Yen  
(74) *Attorney, Agent, or Firm* — Fish & Richardson P.C.

(57)                      **ABSTRACT**

Database query analysis technology, in which an input query is received that relates to multiple attribute classes of data records. An index that is descriptive of the data records in data storage is accessed and relative positions in the accessed index for the multiple attribute classes are identified. The accessed index is processed based on more than one of the multiple attribute classes and the identified relative positions, and the processed index is used to identify groups of data records that are included in the data records in the data storage and that share common values for the multiple attribute classes in the input query. Metrics corresponding to the data records included in each of the identified groups of data records are accessed, computations are performed on the accessed metrics, and a report is generated that reflects results of the computations.

**18 Claims, 19 Drawing Sheets**



(56)

**References Cited**

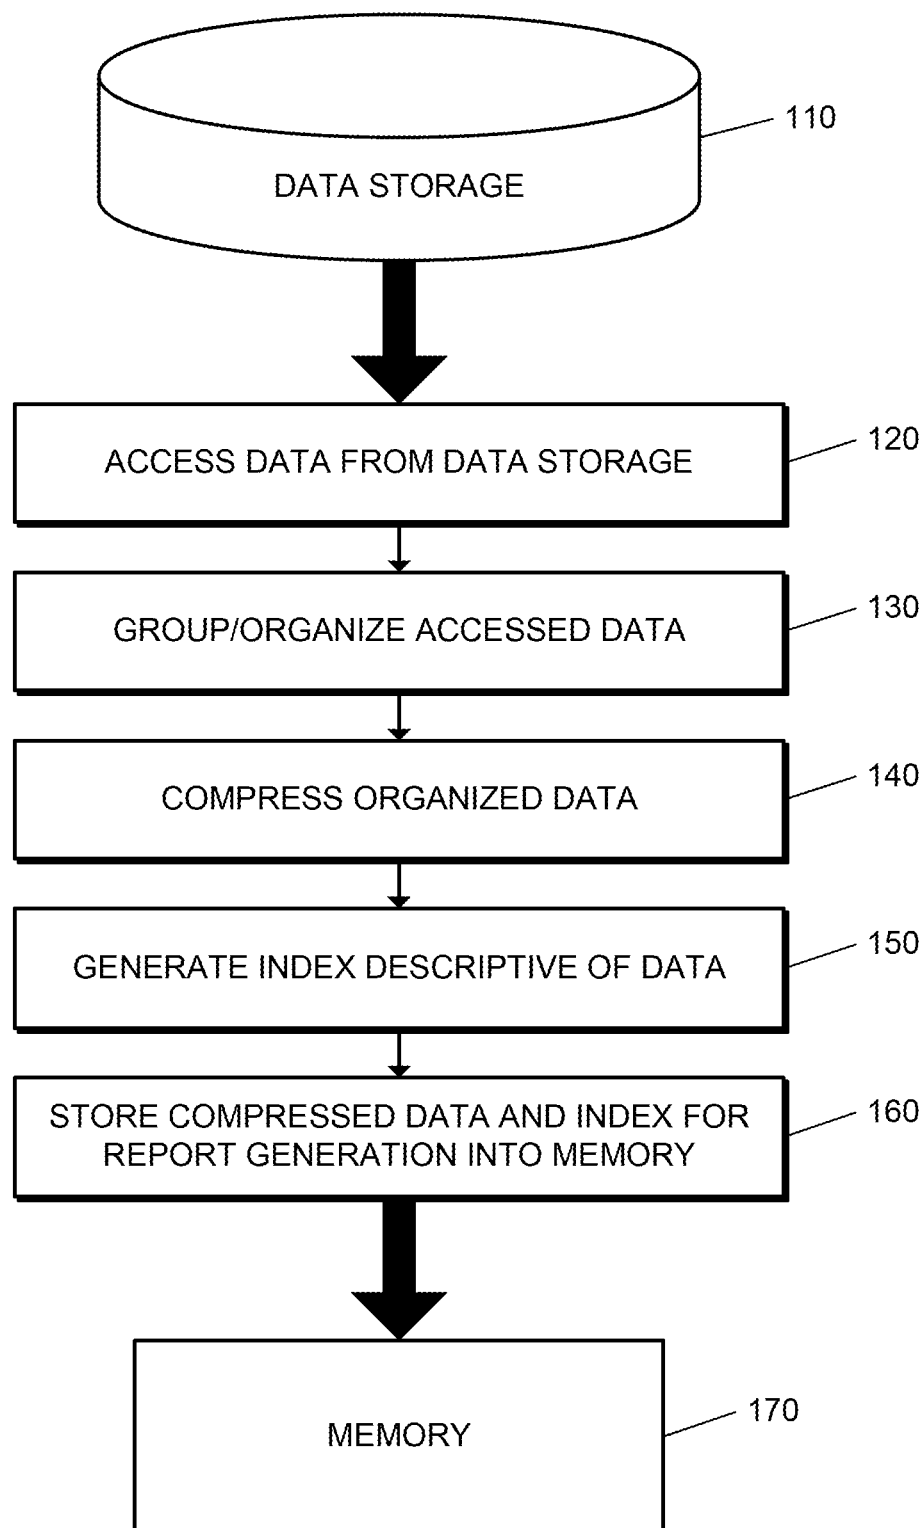
OTHER PUBLICATIONS

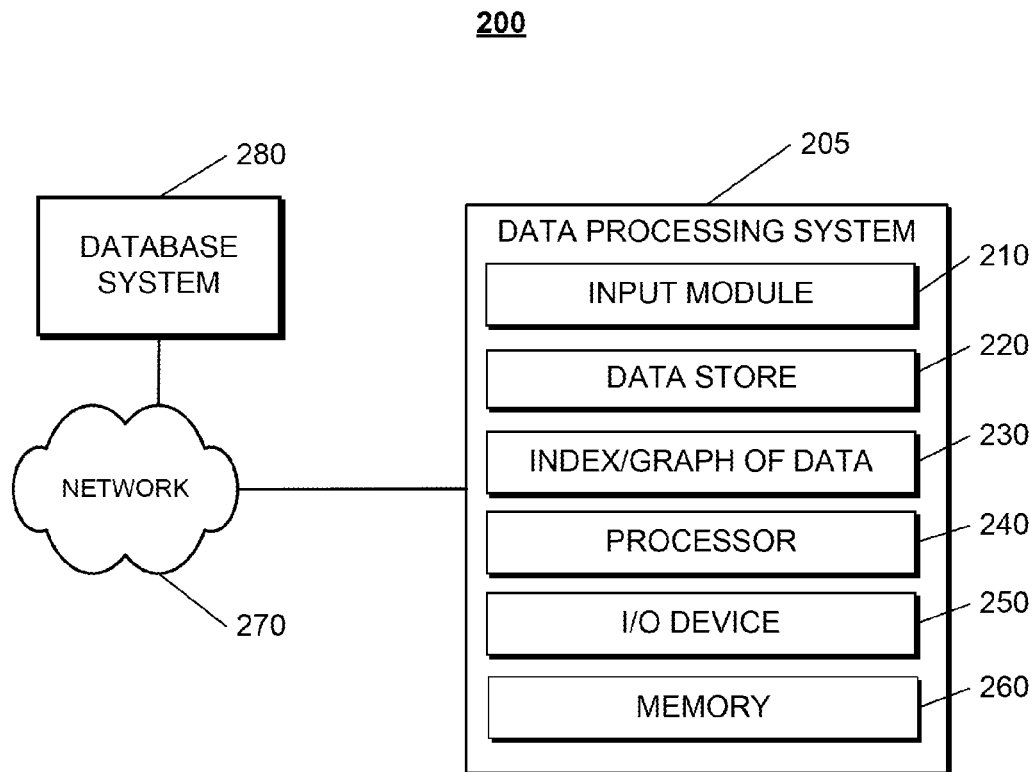
U.S. Final Office Action for U.S. Appl. No. 12/777,631 dated Jan. 3, 2013, 24 pages.

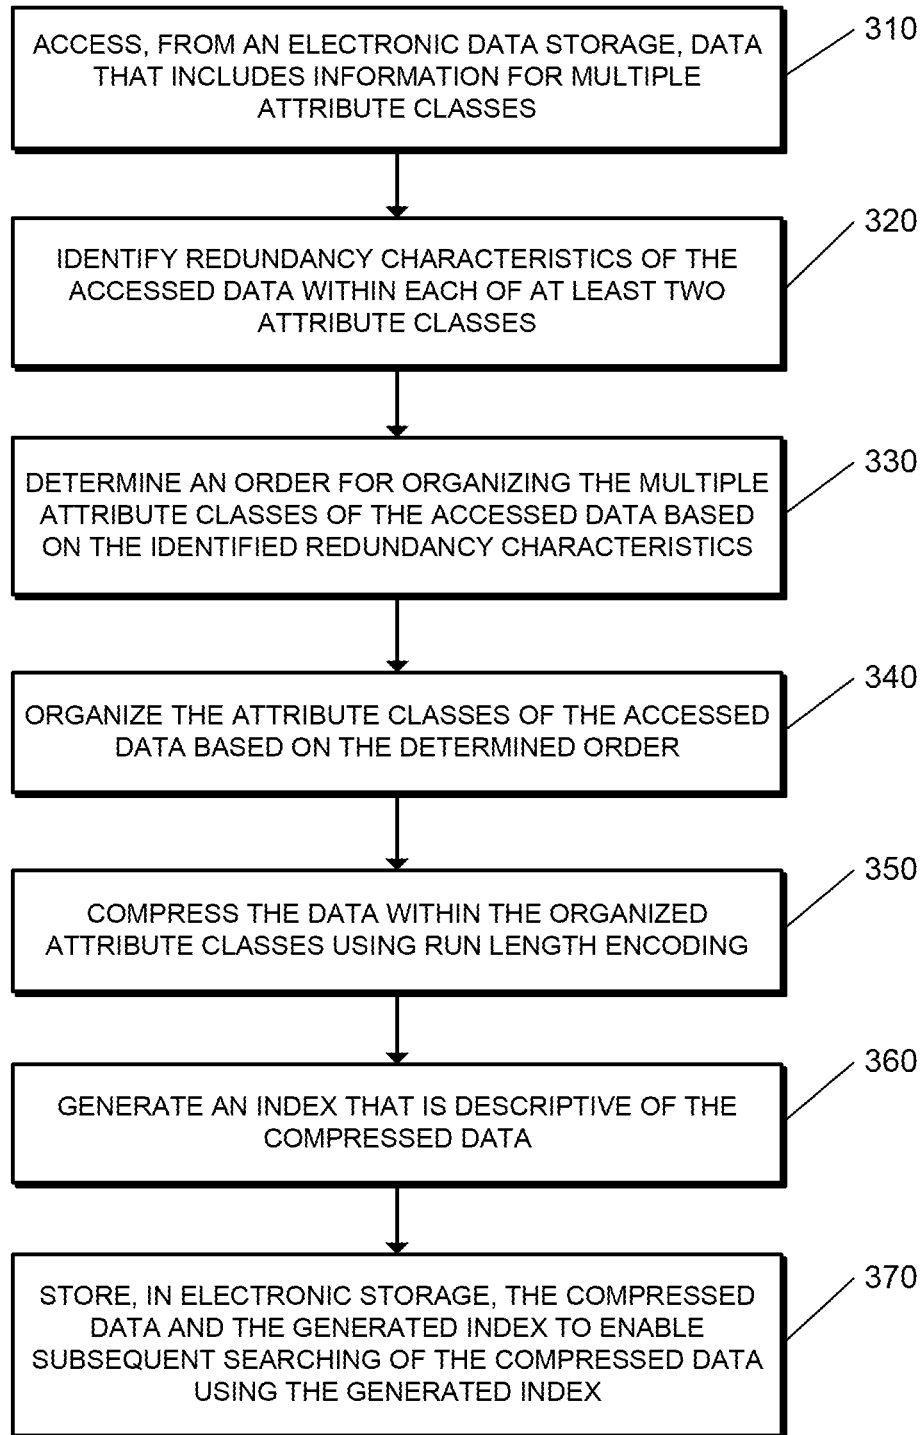
U.S. Notice of Allowance for U.S. Appl. No. 12/777,643 dated May 25, 2012, 5 pages.

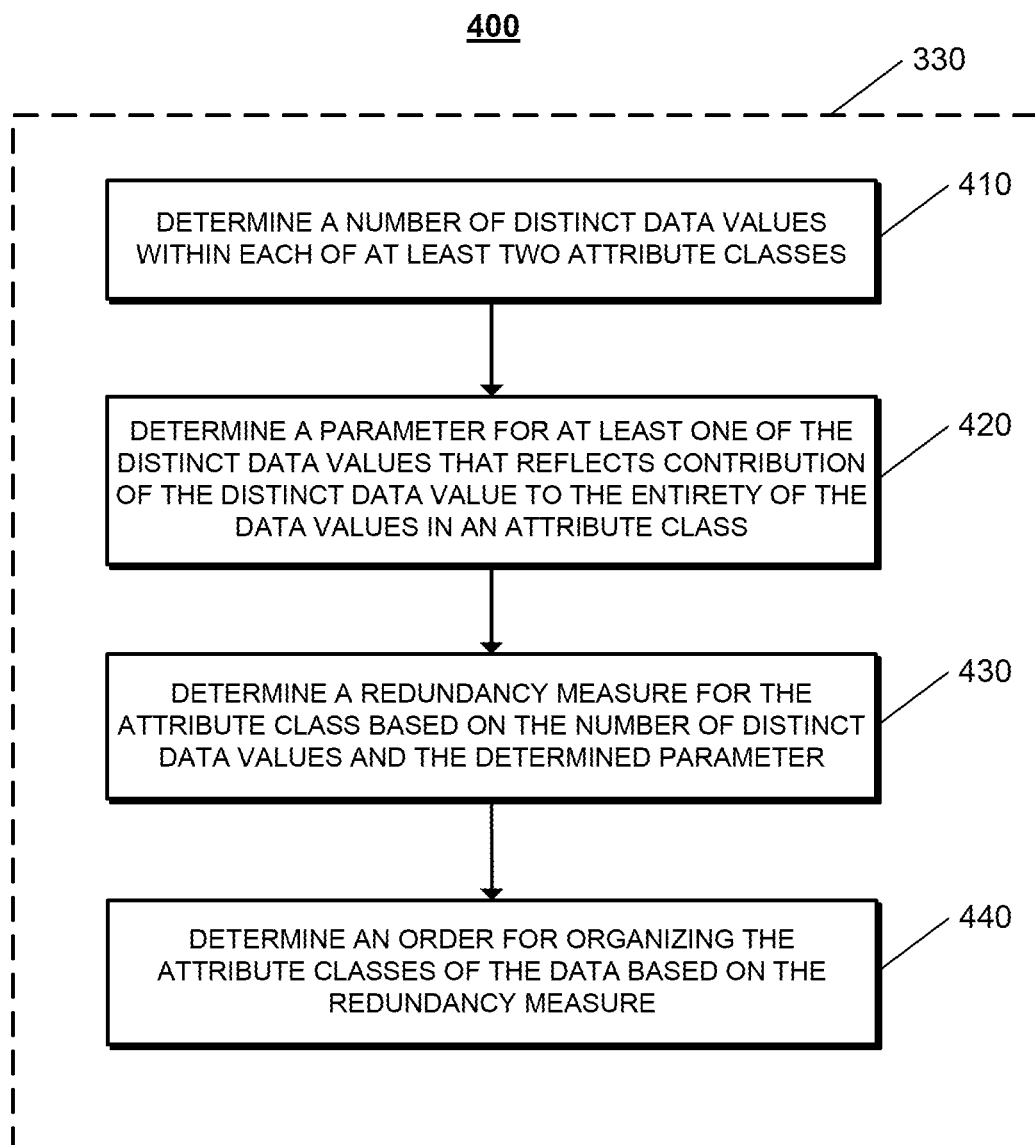
U.S. Notice of Allowance for U.S. Appl. No. 12/777,643 dated Dec. 3, 2012, 12 pages.

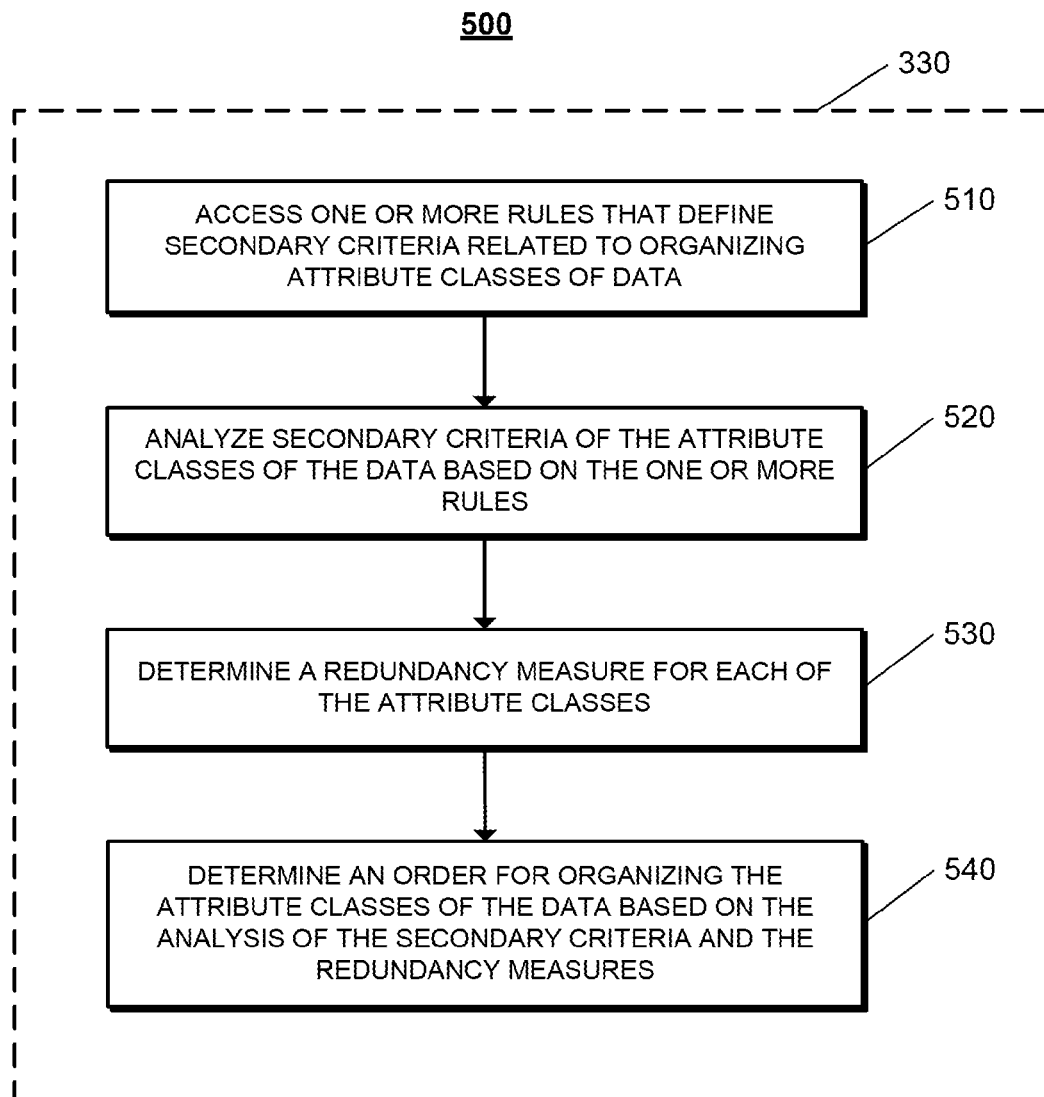
\* cited by examiner

**FIG. 1**

**FIG. 2**

**300****FIG. 3**

**FIG. 4**

**FIG. 5**

610	611	612	613	614	615
	<b>Month</b>	<b>Year</b>	<b>State</b>	<b>City</b>	<b>Sales Metric</b>
621	07/01	07	VA	Vienna	1.2
622	06/02	06	VA	Vienna	1.4
623	06/01	06	VA	McLean	0.8
624	06/02	06	VA	McLean	1.7
625	06/02	06	NY	New York	2.0
626	07/01	07	NY	New York	2.2
627	07/01	07	MD	Baltimore	1.6

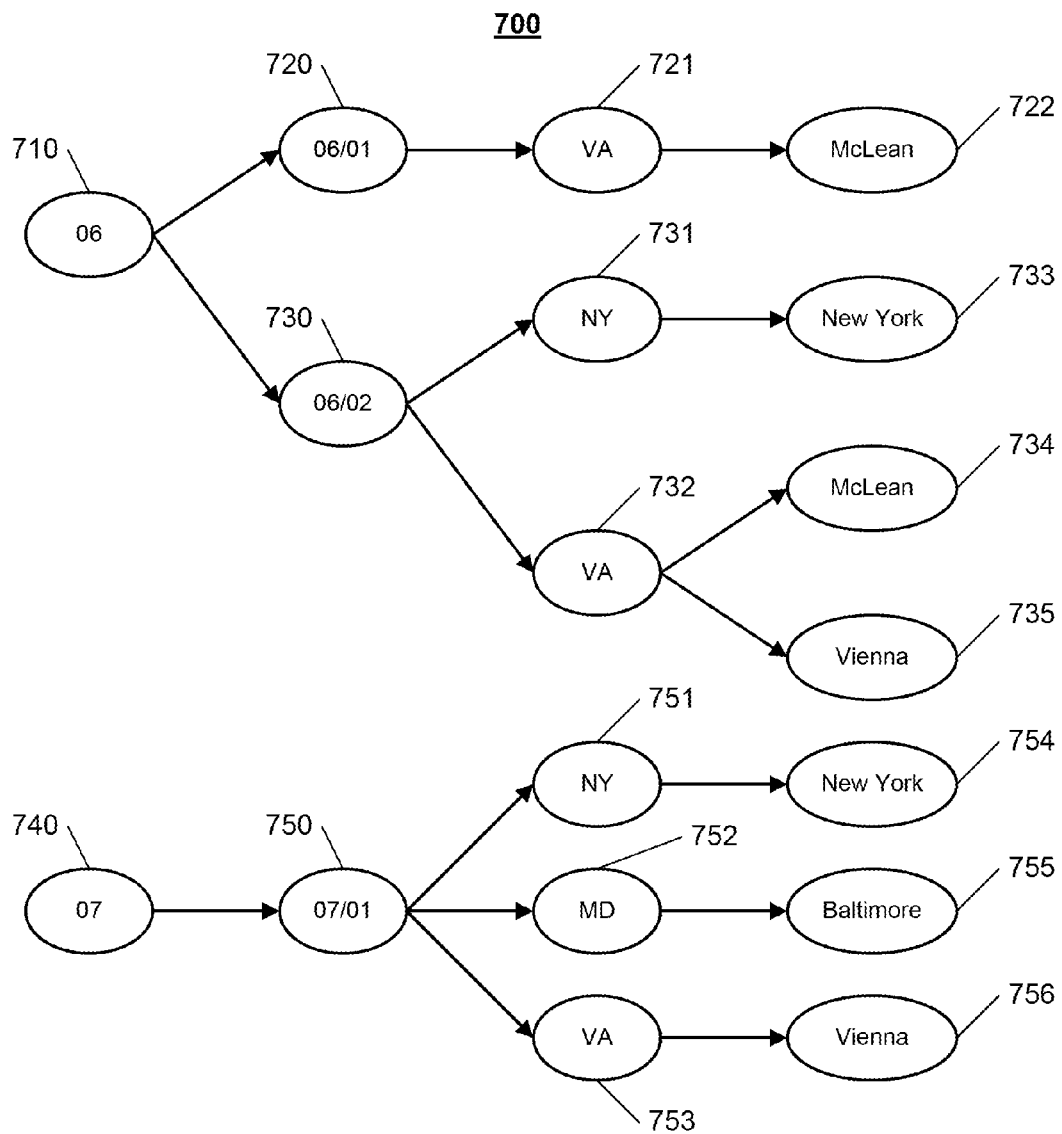
630	612	611	613	614	615
	<b>Year</b>	<b>Month</b>	<b>State</b>	<b>City</b>	<b>Sales Metric</b>
623	06	06/01	VA	McLean	0.8
625	06	06/02	NY	New York	2.0
624	06	06/02	VA	McLean	1.7
622	06	06/02	VA	Vienna	1.4
626	07	07/01	NY	New York	2.2
627	07	07/01	MD	Baltimore	1.6
621	07	07/01	VA	Vienna	1.2

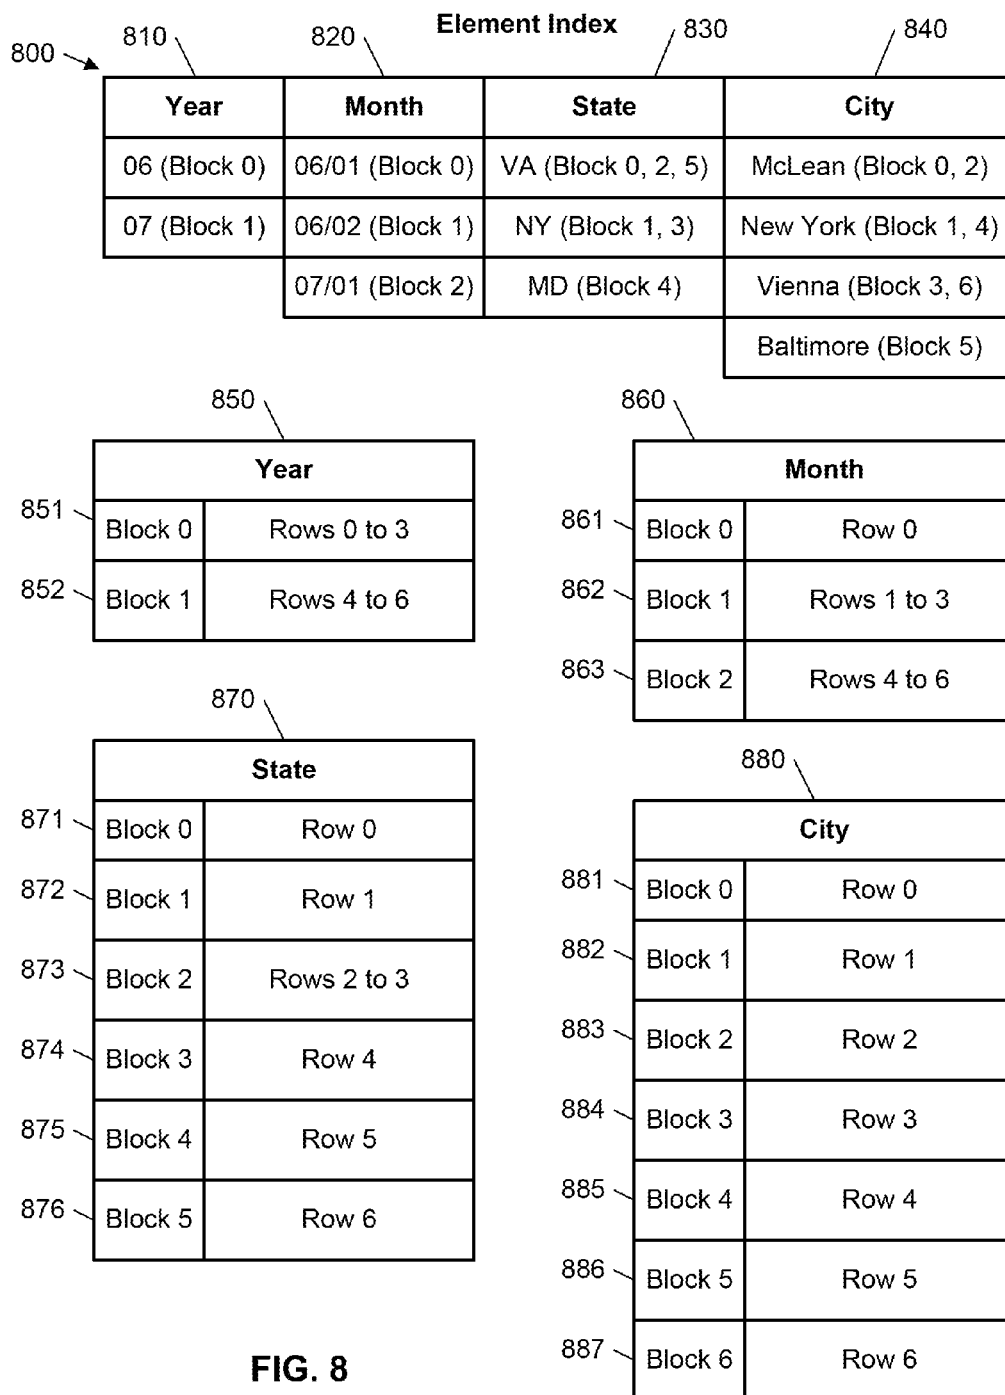
  

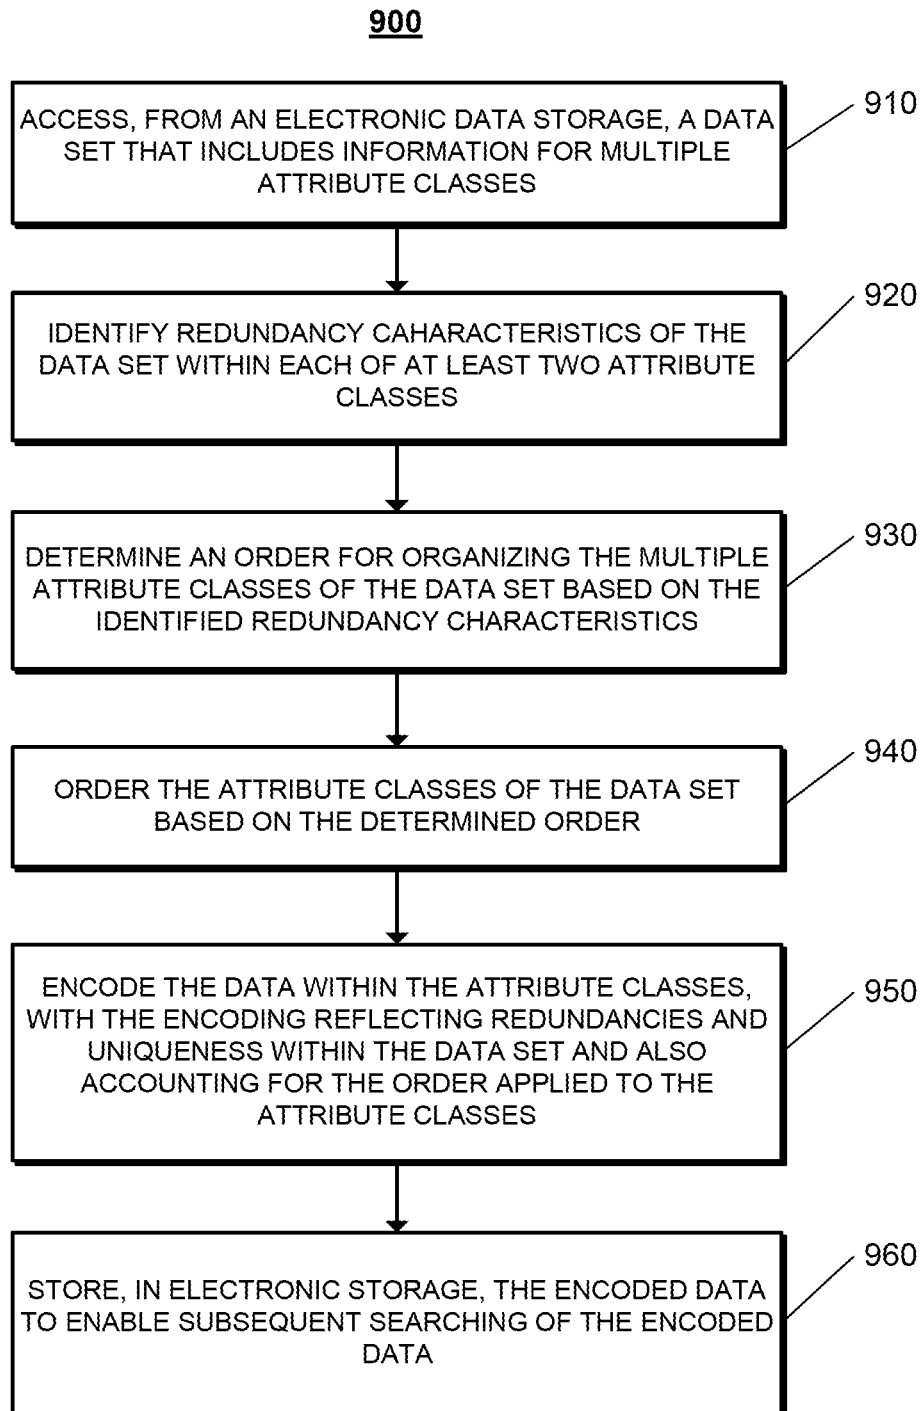
640	612	611	613	614	615
	<b>Year</b>	<b>Month</b>	<b>State</b>	<b>City</b>	<b>Sales Metric</b>
	06	06/01	VA	McLean	0.8
641		06/02	NY	New York	2.0
643			VA	McLean	1.7
644				Vienna	1.4
	07	07/01	NY	New York	2.2
642			MD	Baltimore	1.6
			VA	Vienna	1.2

FIG. 6



**FIG. 7**



**FIG. 9**

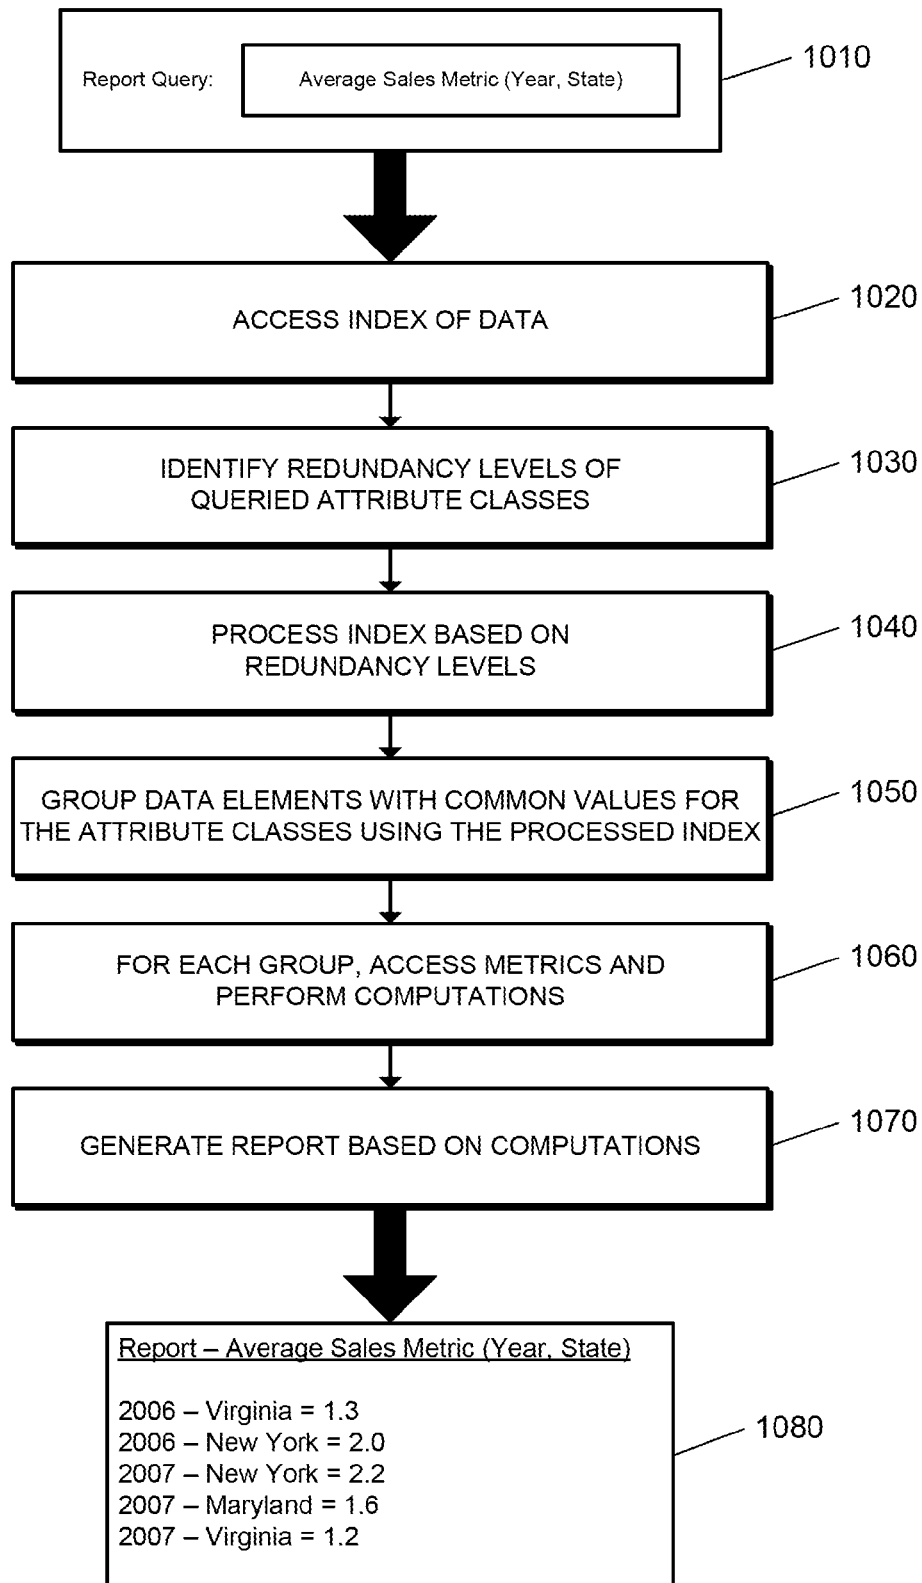
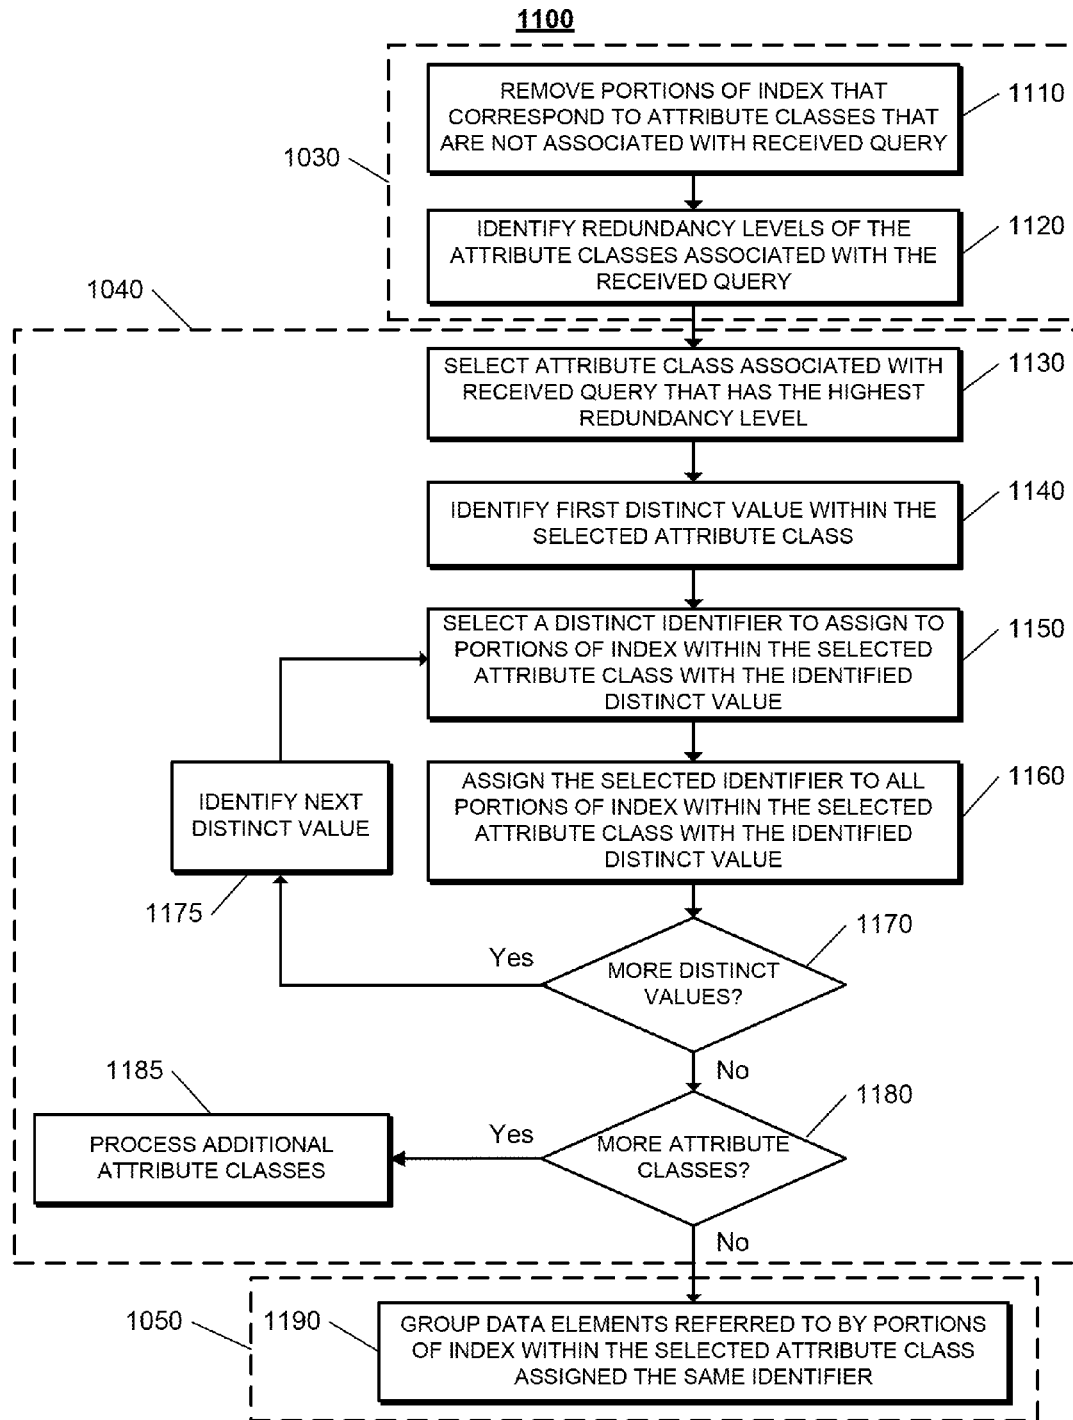
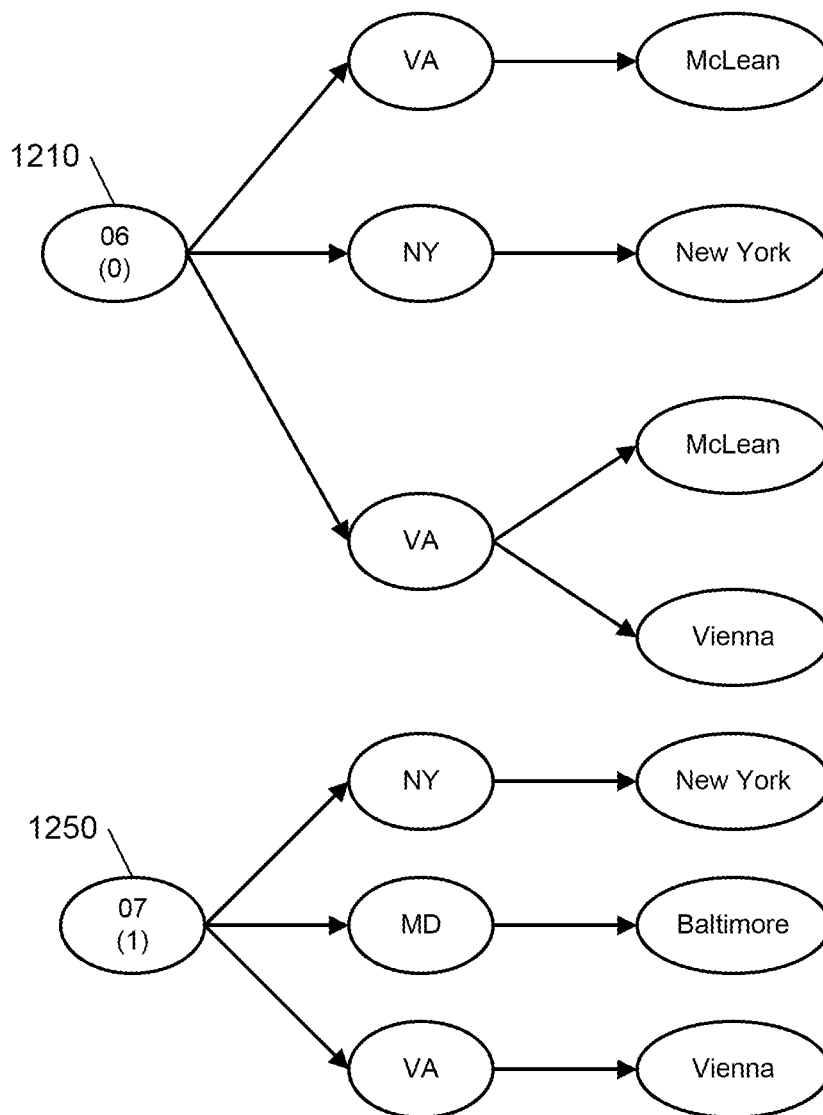


FIG. 10

**FIG. 11**

1200**FIG. 12**

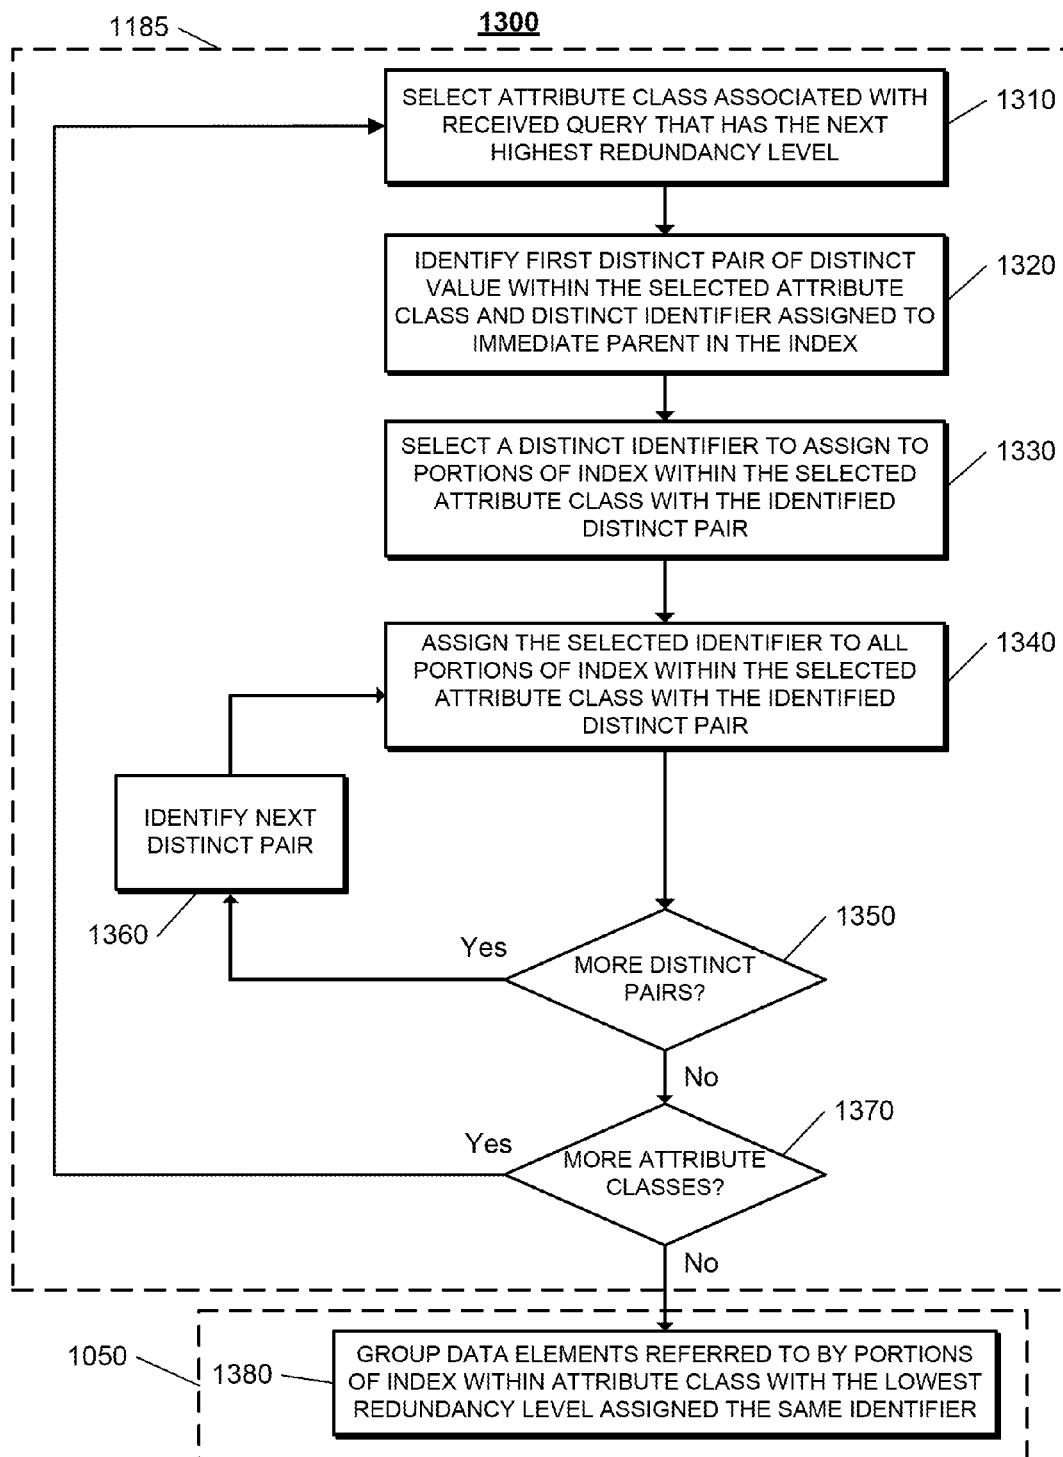
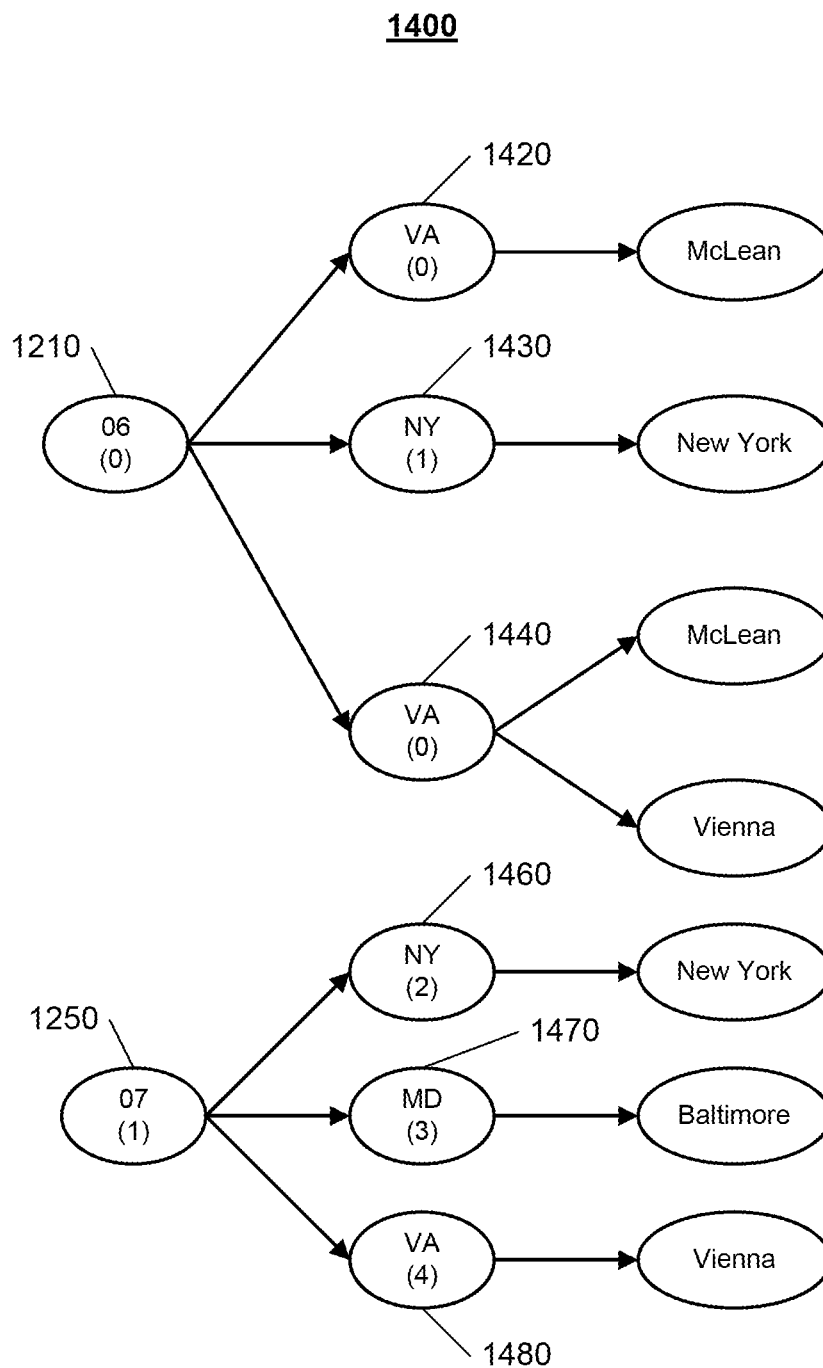
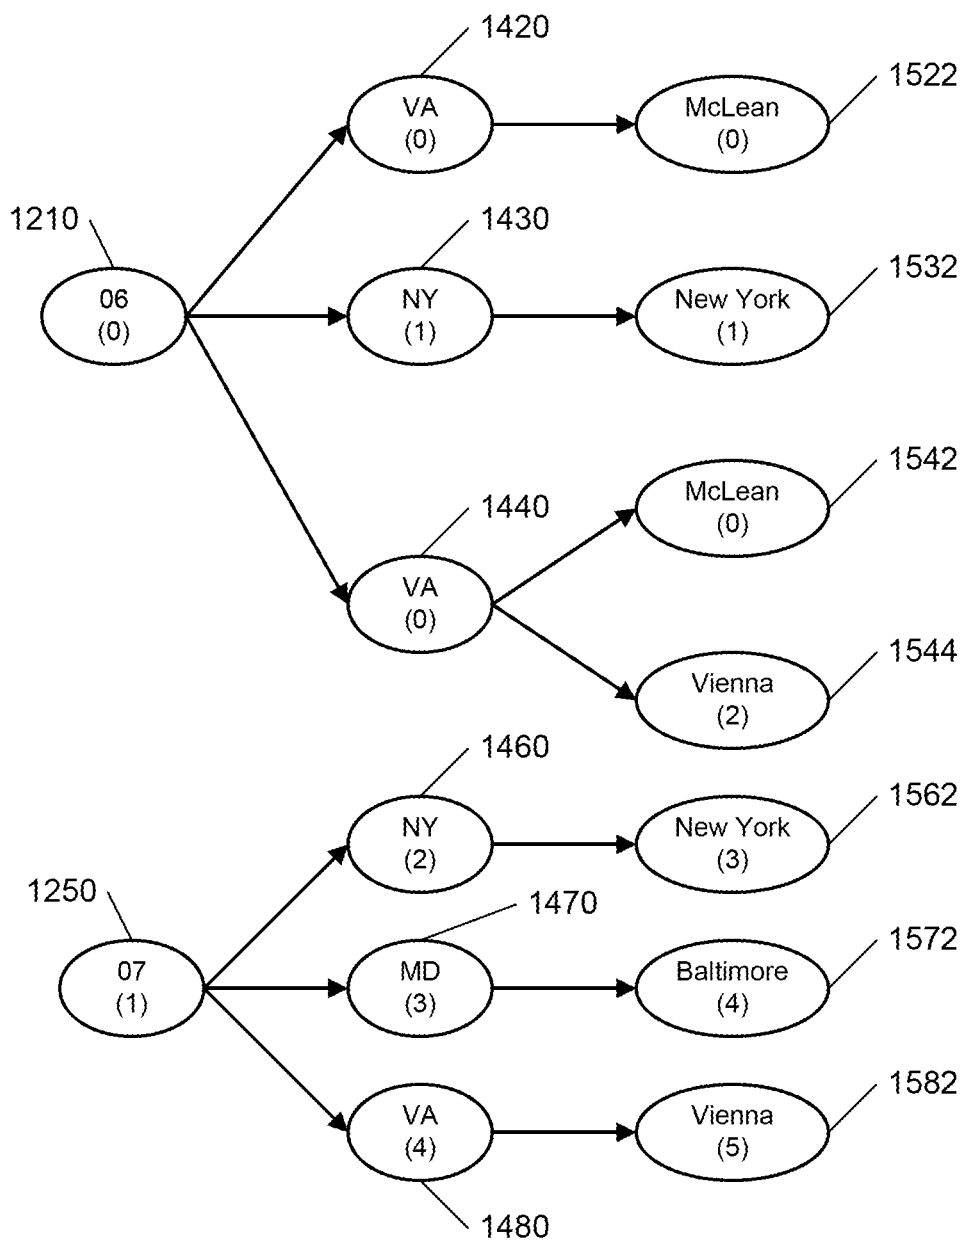
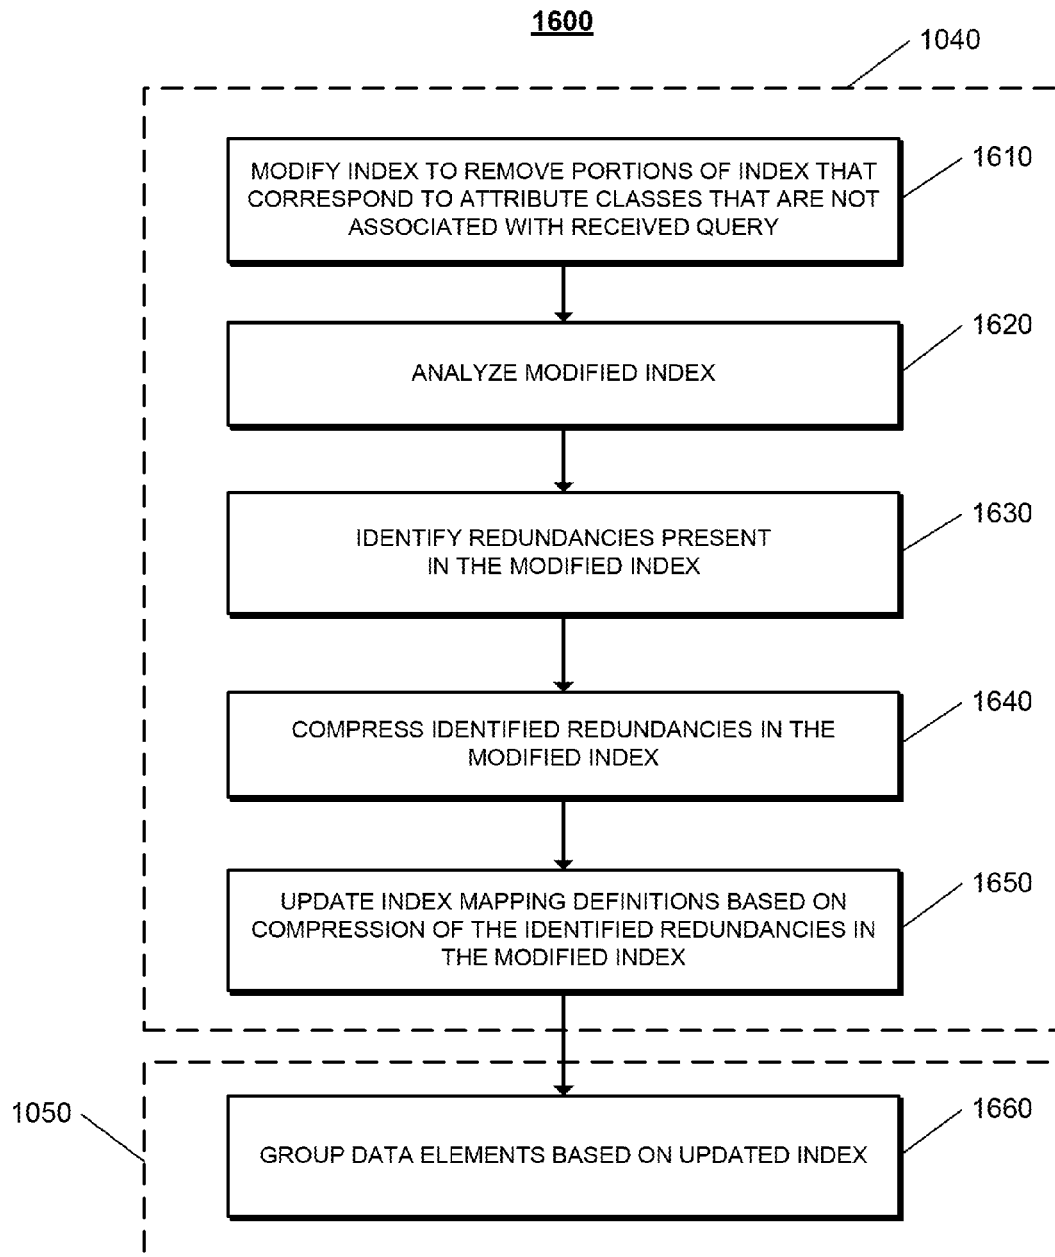


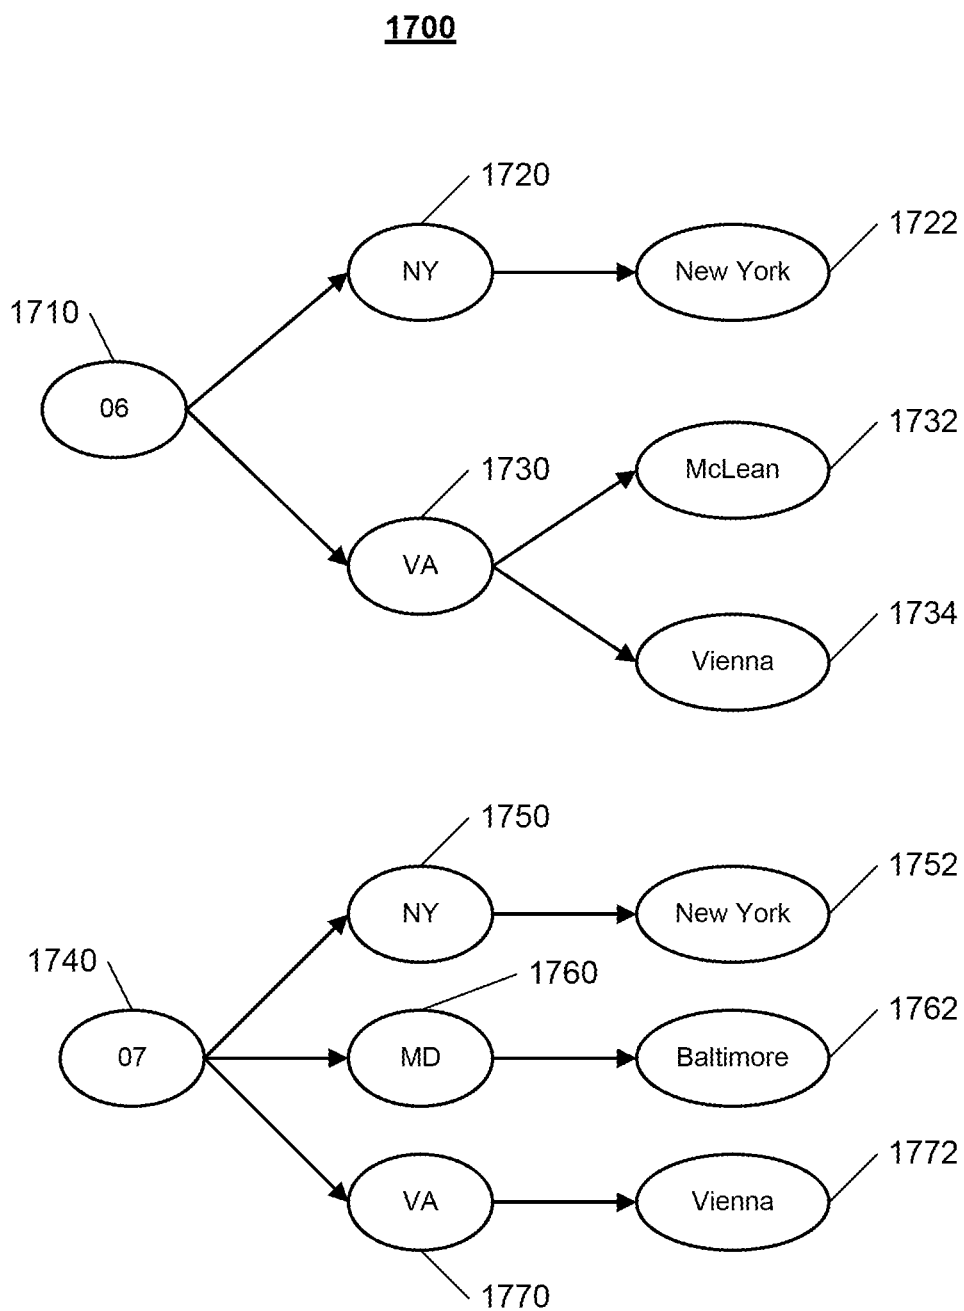
FIG. 13

**FIG. 14**



1500**FIG. 15**

**FIG. 16**

**FIG. 17**

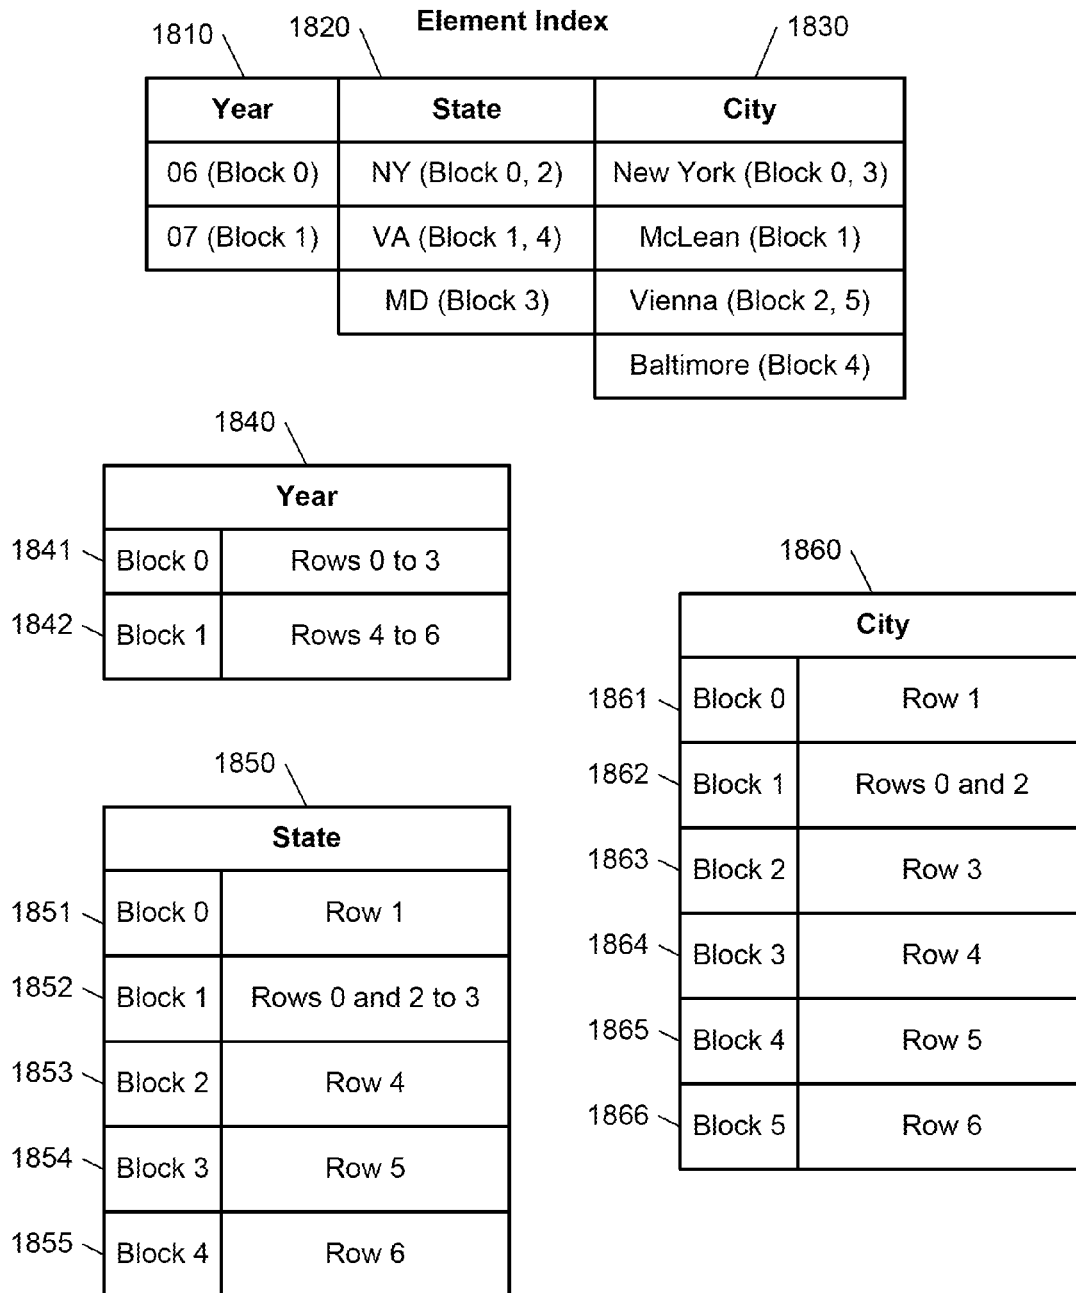


FIG. 18




FIG. 19

1

## INDEX MECHANISM FOR REPORT GENERATION

### CROSS REFERENCE TO RELATED APPLICATIONS

This application is a continuation of U.S. patent application Ser. No. 12/777,643, filed on May 11, 2010, which claims priority to U.S. Provisional Application No. 61/177,489, filed May 12, 2009. The disclosures of the prior applications are considered part of (and are incorporated by reference in) the disclosure of this application.

### TECHNICAL FIELD

This disclosure relates to database query analysis technology.

### BACKGROUND

Computer systems are used to manage and store data. As such, they may be used to analyze data and generate reports based on the analysis results. For instance, computer systems may group and filter data and calculate metric values based on the grouped and filtered data, ultimately providing a report including the calculated metric values.

### SUMMARY

In one aspect, this disclosure relates to database query analysis technology.

Implementations of any of the techniques described throughout the disclosure may include a method or process, a system, or instructions stored on a computer-readable storage device. The details of particular implementations are set forth in the accompanying drawings and description below. Other features will be apparent from the following description, including the drawings, and the claims.

### DESCRIPTION OF DRAWINGS

FIGS. 1, 3-5 and 9-11, 13, and 16 are flowcharts of exemplary processes.

FIGS. 2 and 19 are diagrams of exemplary systems.

FIGS. 6-8, 12, 14, 15, 17, and 18 are diagrams of exemplary data structures.

### DETAILED DESCRIPTION

In some implementations, a system may increase speed and/or efficiency of grouping and/or filtering operations performed in generating service reports for data stored in a data repository (e.g., a database). In these implementations, the system may organize the data stored in the database or data repository in a manner that reduces the time spent in performing grouping and filtering operations. Using the organized data, the time for grouping and filtering operations may be reduced or made negligible as compared to the time in real calculation and, therefore, performance of view report execution may be improved.

For example, the system may represent attribute information of the data using blocks. The system may first re-order and sort the attribute information and compress the data into blocks using run-length encoding. The system then may build an index on the blocks, and may use the blocks and index structure to perform filtering and grouping operations in a relatively efficient manner.

2

FIG. 1 illustrates a process for organizing data and generating an index of the organized data for report generation. The process shown in FIG. 1 is described generally as being performed by a processor. In some implementations, the process shown in FIG. 1 may be performed by one or more processors included in one or more electronic devices or may be performed any type of electronic device (e.g., a server, a computer, etc.).

The processor accesses data from a data storage 110 (120). For instance, the processor may retrieve data from a database or data warehouse using a data access command or a query (e.g., an SQL query statement). The processor may access data over a direct connection to the data storage 110 or over a network.

In some examples, the processor defines relationships between data attribute classes for the data stored in the data storage 110 based on input provided by a database architect. For instance, based on input from the architect, the processor may structure the data to be stored as cubes that have dimensions defining relationships between attribute classes. Each dimension may define related attribute classes in a parent-child relationship (e.g., a dimension of time may have a parent class of "year" and child classes of "month," "day," and "hour"). The processor may determine what data to access from the data storage 110 based the type of data queried by users (e.g., data of dimensions of interest as identified by a client) or based on input provided by a database architect. The processor may access the data identified by the database architect as being of interest without accessing all of the data stored in the data storage 110.

The processor groups or organizes the accessed data (130). The processor may process the data accessed from the data storage 110 and group or organize the data in a manner that is more efficient than the manner in which the data is stored in the data storage 110. For example, the processor may rearrange relationships of attribute classes in the data differently than the data storage 110. In this example, the processor may arrange columns of data in a table differently than a table stored in the data storage 110. The processor also may sort the accessed data to group similar (or the same) data values together within the accessed data. The processor further may filter any of the accessed data that is not needed for report generation.

In some implementations, the processor may determine how to group or organize the accessed data by identifying redundancies within the accessed data and organizing the data in a manner that leverages the identified redundancies. By leveraging the identified redundancies, the processor may be able to reduce the storage capacity needed to store the accessed data and also may be able to reduce the processing time needed to locate relevant portions of the accessed data. Grouping and organizing accessed data based on identified redundancies is described in more detail below with respect to FIGS. 3-5.

The processor compresses the organized data (140). Based on the organization of the data, the processor may compress the data to reduce the storage size of the data and reduce the number of operations needed to be performed to identify relevant portions of the data (e.g., reduce the number of comparisons needed to be made to execute a query). The processor may compress data within particular attribute classes using run length encoding to generate a set of blocks of the same data values for the corresponding attribute class. The blocks may require less storage capacity than the accessed data and evaluating operations for the blocks may be more efficient than evaluating operations on individual records within the accessed data.

The processor generates an index that is descriptive of the organized data (150). The index may be a data structure that defines the organization of the data. For instance, when the data is organized and compressed into blocks, the index may identify blocks within the organized data and identify relationships between the blocks in the organized data. The relationships may indicate whether blocks are within the same attribute class and whether a particular block is related to other blocks within other attribute classes (e.g., whether the block has a parent block associated with a parent attribute class and/or whether the block has a child block associated with a child attribute class). The processor may generate the index by identifying the blocks within the organized and compressed data, determining relationships between the identified blocks, and generating data that is descriptive of the identified blocks and determined relationships.

The processor also may store addressing information (e.g., row numbers) for each of the identified blocks to enable future accesses of data stored within identified blocks by referencing the index. The processor may use the index to improve the speed of operations on the accessed data by providing efficient access of ordered records.

The processor stores the compressed data and the index for report generation into a memory 170 (160). For instance, the processor may store the compressed data and the index into any type of random access memory. Storing the compressed data in the memory 170 may enable faster report generation because the access time of accessing data from the memory 170 may be faster than the access time of accessing data from the data storage 110. When performing a report generation process, the processor may access the index stored in the memory 170 and use the index to identify locations in the memory 170 for the relevant portions of the compressed data. The processor may access data from the identified locations in the memory 170 and use the accessed data to generate a report.

Referring to FIG. 2, a block diagram of a system 200 is shown. The system 200 includes a data processing system 205, a network 270, and a database system 280. The network 270 enables the data processing system 205 and the database system 280 to exchange electronic communications.

The data processing system 205 includes an input module 210, a data store 220, index or graph data 230, a processor 240, an input/output (I/O) device 250, and a memory 260. The data processing system 205 may be used to satisfy queries and generate reports based on data stored in the database system 280. The data processing system 205 may be a general purpose computer, server, or any other type of electronic device that includes electronic components that are capable of accessing and processing data. The data processing system 205 may be implemented within hardware or a combination of hardware and software.

The input module 210 imports data associated with a report generation process. The data may include data from a database that is used to generate a report (e.g., data from a business database or transaction processing system). The input module 210 may input data from a device (e.g., the database system 280) connected to the network 270. In some implementations, the input module 210 reformats and/or transforms the data such that the data may be processed and stored by other components within the data processing system 205.

The data processing system 205 also includes a data store 220. In some implementations, data from the input module 210 is stored in the data store 220. The data store 220 may be, for example, a database that logically organizes data into a series of database tables. The data store 220 may be a hard

disk drive, non-volatile memory (e.g., Flash memory), or another type of electronic storage device.

The data processing system 205 also includes index or graph data 230. The index or graph data 230 may include a data structure that defines the organization of data that is processed in satisfaction of a report generation command. The data structure may identify relationships within the data and include addressing information that maps portions of the index or graph to actual storage locations where the data resides. In some implementations, the index or graph data 230 may be received, by the data processing system 205, from the database system 280.

The data processing system 205 also includes a processor 240. The processor 240 may be a processor suitable for the execution of a computer program such as a general or special purpose microprocessor, and any one or more processors of any kind of digital computer. Generally, a processor receives instructions and data from a read-only memory or a random access memory or both. The processor 240 receives instructions and data from the components of the data processing system 205 to, for example, organize and compress data and generate the index or graph data 230. The processor 240 also may receive instructions and data from the components of the data processing system 205 to generate a report in satisfaction of a query using the index or graph data 230. In some implementations, the data processing system 205 includes more than one processor.

The data processing system 205 also includes the I/O device 250, which is configured to allow user input. For example, the I/O device 250 may be a mouse, a keyboard, a stylus, a touch screen, a track ball, a toggle control, one or more user input buttons, a microphone, or any other device that allows a user to input data into the data processing system 205 or otherwise communicate with the data processing system 205. The I/O device 250 may receive input from a user that defines a query or a report generation command. In some implementations, the user may be a machine and the user input may be received from an automated process running on the machine. In other implementations, the user may be a person.

The I/O device 250 also may include a device configured to output generated reports and status information. For instance, the I/O device 250 may include a display device configured to display generated reports and status information. The I/O device 250 also may include a speaker configured to provide audible output.

The data processing system 205 also includes a memory 260. The memory 260 may be any type of tangible machine-readable storage medium. The memory 260 may, for example, store the data included in the data store 220 and/or the index or graph data 230. In some implementations, the memory 260 may store instructions that, when executed, cause the data processing system 205 to, for example, organize and compress data and generate the index or graph data 230.

The system 200 also includes a network 270. The network 270 is configured to enable exchange of electronic communications between devices connected to the network 270. For example, the network 270 may be configured to enable exchange of electronic communications between the data processing system 205 and the database system 280. The network 270 may include, for example, one or more of the Internet, Wide Area Networks (WANs), Local Area Networks (LANs), analog or digital wired and wireless telephone networks (e.g., a PSTN, Integrated Services Digital Network (ISDN), a cellular network, and Digital Subscriber Line (DSL)), radio, television, cable, satellite, or any other deliv-

ery or tunneling mechanism for carrying data. Network 270 may include multiple networks or subnetworks, each of which may include, for example, a wired or wireless data pathway. The network 270 may include a circuit-switched network, a packet-switched data network, or any other network able to carry electronic communications. For example, the network 270 may include networks based on the Internet protocol (IP) or asynchronous transfer mode (ATM).

The database system 280 is an electronic device configured to store data and exchange communications with the data processing system 205 (e.g., multiple data processing systems) over the network 270. For example, the database system 280 may be configured to store an organization's data and output the organization's data in response to requests (e.g., SQL statements or queries). In this example, the database system 280 may exchange communications with the data processing system 205 to receive input defining data needed from the database system 280 and provide the data needed as output to the data processing system 205. The database system 280 may include one or more databases and/or data warehouses.

Although the example data processing system 205 is shown as a single integrated component, one or more of the modules and applications included in the data processing system 205 may be implemented separately from the data processing system 205 but in communication with the data processing system 205. For example, the data store 220 may be implemented on a centralized server that communicates and exchanges data with the data processing system 205. In this example, the database system 280 may communicate with the data processing system 205 and perform operations described above as being performed by the data processing system 205 or may perform operations that assist the data processing system 205 performing operations described throughout the disclosure.

FIG. 3 illustrates a process 300 for organizing and compressing data and generating an index to enable subsequent searching of the organized and compressed data using the generated index. The operations of the process 300 are described generally as being performed by the system 200. The operations of the process 300 may be performed exclusively by the data processing system 205, may be performed exclusively by the database system 280, or may be performed by a combination of the data processing system 205 and the database system 280. In some implementations, operations of the process 300 may be performed by one or more processors included in one or more electronic devices.

The system 200 accesses, from an electronic data storage, data that includes information for multiple attribute classes (310). The system 200 may retrieve data from a database or data warehouse using a data access command or a query (e.g., an SQL query statement). For instance, the data processing system 205 may send, over the network 270, a data access request to the database system 280 and the database system 280 may send, over the network, the requested data to the data processing system 205.

Although the accessed data may not include data for all of the attribute classes for the data stored in the electronic data storage (although it may), the accessed data includes information for multiple attribute classes. For example, the accessed data may be stored as cubes that have dimensions defining relationships between attribute classes. Each dimension may define related attribute classes in a parent-child relationship (e.g., a dimension of time may have a parent class of "year" and child classes of "month," "day," and "time"). The system 200 may access a cube of data that includes one or more dimensions that define a relationship between multiple

attribute classes. The system 200 also may access multiple columns worth of data from a database table.

The system 200 may determine which data to access based on rules defined by a database architect or system administrator. For example, a database architect or system administrator may set rules defining data of interest to an organization. In this example, the rules may define which attribute classes are of interest to an organization and the system 200 accesses the data for the attribute classes of interest. The rules also may define time periods of interest to an organization and the system 200 may access data associated with the relevant time periods (e.g., data stored within the last five years).

In some implementations, the system 200 may determine which data to access dynamically based on the user or device requesting access. In these implementations, the system 200 may determine access level credentials of the user or device requesting access to the data and determine which data to access based on the determined credentials. In addition, the rules may define that different users or different types of users receive different attribute classes of data. For instance, the system 200 may access financial data for an organization when the user accessing the data is a financial analyst, but may access personnel data for the organization when the user accessing the data is a human resources manager.

In some examples, the system 200 may access data from the electronic data storage prior to receiving a report generation command such that the data is pre-loaded for execution of a report generation process. In these examples, the system 200 may access the data when a user logs onto the system 200 or when the system 200 is powered on. The system 200 also may access data at periodic intervals, such as one time each day.

The system 200 identifies redundancy characteristics of the accessed data within each of at least two attribute classes (320). The system 200 may identify a number of distinct values within each of the at least two attribute classes as the redundancy characteristics. For example, the system 200 may process the accessed data by analyzing each data value for an attribute class and counting the number of distinct values that exist for the attribute class in the stored data. In this example, the system 200 may sequentially analyze all of the records in the accessed data, track data values for the attribute class present in the data records (e.g., store analyzed values in temporary storage), and compare data values for subsequent records to the tracked values to determine whether the data values are distinct from other data values included in the data records. When a data value matches a tracked value, the system 200 determines that the data value is not distinct (e.g., determines that the data value is redundant of at least one other data value) and continues processing the next data record without updating tracked data. When a data value does not match any tracked value, the system 200 determines that the data value is distinct (e.g., determines that the data value is not redundant of at least one other data value), stores the data value with the tracked data values for comparison against subsequent records, and increments a counter that tracks the number of distinct data values within the attribute class.

In some implementations, the system 200 may sort the accessed data with respect to an attribute class of interest prior to identifying the number of distinct values within the attribute class of interest. Sorting the accessed data may improve efficiency in identifying the number of distinct values because the data records with the same data value for the attribute class would be arranged together and processed consecutively. Accordingly, because the system 200 knows the data values are arranged consecutively, the system 200 may only have to compare a data value to the most recently tracked



data value. Specifically, if the data value is redundant of a previously processed data value, it is necessarily redundant of the most recently tracked data value because it would have been grouped together with the most recently tracked data value in the sorting process.

The system **200** may calculate other measures of data redundancy within an attribute class to identify redundancy characteristics. For example, the system **200** may determine a percentage, within each of the at least two of the multiple attribute classes, of the accessed data that has a redundant value for the corresponding attribute class. In this example, the system **200** may compute the percentage as the number of distinct data values over the total number of data values.

The system **200** also may determine a distribution of redundant values within an attribute class. For instance, for each distinct data value within an attribute class, the system **200** may determine the number or percentage of records that include the distinct data value. The system **200** may use the distribution of redundant values to determine the benefit of leveraging the redundancy of the data within the attribute class. For example, a first attribute class may have the same number of distinct data values as a second attribute class, but a single, distinct data value within the first attribute class may be present in a relatively high percentage of the data records while the distinct data values in the second attribute class may be more evenly distributed. In this example, the system **200** may determine characteristics of the distribution of redundant data values within the first attribute class and the second attribute class and determine a metric that corresponds to the degree with which the redundancy of the data may be leveraged in compressing the data. The metric corresponding the first attribute class may reflect a higher degree of being able to leverage redundancy of the data than the metric corresponding the second attribute class because the relatively high degree of redundancy of the single, distinct data value in the first attribute class may be leveraged more so than any of the redundant data values in the second attribute class.

In some examples, the system **200** may consider redundancy of data values within related (e.g., child) attribute classes as part of the redundancy characteristics. In these examples, data within parent and child attribute classes may need to be stored together (e.g., when the data for the parent attribute class and the child attribute class is stored in a single record) and, therefore, redundancy characteristics of the data within the child attribute class may impact the ability to leverage redundancy within the parent attribute class. The system **200** may group distinct data values in a parent attribute class together and, for each group within the parent attribute class (e.g., each block of the same data value within the parent attribute class), the system **200** may determine the number of distinct data values in a child attribute class that are associated with the corresponding group. Accordingly, rather than analyzing redundancy of the child attribute class as a whole, redundancy of the child attribute class is measured based on groups of redundant data within the parent attribute class (e.g., two of the same value in the child attribute class may be counted as distinct when the two values are associated with different groups in the parent attribute class). In this regard, the system **200** may measure a level of redundancy in a combination of the parent and child attribute classes. This may provide a measure of the ability of further leverage the redundancy of the parent attribute class within the child attribute class.

For example, a first parent attribute class may have a greater number of distinct data values than a second parent attribute class such that, taken alone, the second parent attribute class has a greater level of data redundancy than the

first parent attribute class. However, in this example, the first parent attribute class may be associated with a first child attribute class that has a relatively high level of data redundancy for each group of distinct data values in the first parent attribute class when data records are grouped into blocks of distinct values in the first parent attribute class (as an extreme example, suppose the first child attribute class has a single distinct data value for each group). The second parent attribute class may be associated with a second child attribute class that has a relatively low level of data redundancy for each group of distinct data values in the second parent attribute class when data records are grouped into blocks of distinct values in the second parent attribute class (as an extreme example, suppose the second child attribute class has a distinct data value for each data record included in each group). In this example, because of the child attribute classes, the redundancy characteristics of the first attribute class may be leveraged better than the redundancy characteristics of the second attribute class, even though, taken alone, the second parent attribute class has a greater level of data redundancy than the first parent attribute class. The system **200** may identify the redundancy within the child attribute classes and track data that reflects combined redundancy as part of the identified redundancy characteristics.

The system **200** determines an order for organizing the multiple attribute classes of the accessed data based on the identified redundancy characteristics (**330**). For instance, when the system **200** identifies a number of distinct values within each of the multiple attribute classes, the system **200** may determine a relative order among the multiple attribute classes of the accessed data based on the identified number of distinct values within each of the multiple attribute classes. In this regard, the system **200** may determine to order the multiple attribute classes by ordering attribute classes with a lower number of distinct data values prior to attribute classes with a higher number of distinct data values. The system **200** may determine that a first attribute class has a lower number of distinct values than a second attribute class and, therefore, order the first attribute class prior to the second attribute class in the determined relative order.

In some examples, the system **200** may identify, from among the multiple attribute classes, an attribute class having a lowest number of distinct values and order the identified attribute class having the lowest number of distinct values first in the determined relative order. In these examples, the system **200** may order the remaining attribute classes by increasing number of distinct data values.

When the system **200** determines a percentage, within each of the multiple attribute classes, of the accessed data that has a redundant value for the corresponding attribute class, the system **200** may determine a relative order among the multiple attribute classes of the accessed data based on the determined percentages. For instance, the system **200** may identify the attribute class that has the highest percentage and order the identified attribute class having the highest percentage first in the determined relative order. Also, the system **200** may order the attribute classes in an order of decreasing percentages.

In some implementations, after the system **200** identifies the attribute class to order first in the determined relative order (e.g., the attribute class with the lowest number of distinct data values, the attribute class with the highest percentage of redundant data, etc.), the system **200** may reevaluate data redundancy characteristics of the remaining attribute classes based on the determination of the first attribute class. In these implementations, the data redundancy characteristics of the other attribute classes may change based on which attribute

class is determined to be first in the order. For example, after identifying the first attribute class, the system 200 may organize the data and identify blocks of redundant data included in the first attribute class. In this example, the system 200 may identify data redundancy characteristics of data in the other attribute classes within the blocks of redundant data identified in the first attribute class. Because data that is otherwise redundant in the other attribute classes, may span multiple, different blocks of redundant data in the first attribute class, the system 200 may not be able to fully leverage the redundancy of the data and the redundancy characteristics may change. As such, within the other attribute classes, a second attribute class may have redundancy characteristics that reflect a higher degree of data redundancy than redundancy characteristics of a third attribute class prior to the selection of the first attribute class in the order. After selection of the first attribute class in the order, however, the second attribute class may have updated redundancy characteristics that reflect a lower degree of data redundancy than updated redundancy characteristics of the third attribute class. Specifically, the distribution of redundant data within the third attribute class may be relatively similar to the distribution of redundant data within the first attribute class and the distribution of redundant data within the second attribute class may be relatively dissimilar to the distribution of redundant data within the first attribute class. Therefore, the third attribute class may have a higher degree of data redundancy than the second attribute class when data redundancy characteristics are determined after establishing the first attribute class in the order.

The system 200 may determine updated redundancy characteristics for each of the remaining attribute classes based on the selection of the first attribute class and identify a next attribute class in the order based on the updated data redundancy characteristics. For instance, the system 200 may select the second attribute class in the order as the remaining attribute class that has updated data redundancy characteristics that reflect the highest degree of data redundancy based on selection of the first attribute class in the order. After selection of each attribute class in the order, the system 200 may continue to update data redundancy characteristics for the remaining attribute classes and identify the next attribute class as the attribute class having updated data redundancy characteristics that reflect the highest degree of data redundancy in light of the prior selections.

FIG. 4 illustrates a process 400 for determining an order for organizing multiple attribute classes of data based on identified redundancy characteristics. The process 400 may be used in determining an order for organizing multiple attribute classes of data based on identified redundancy characteristics referenced above with respect to reference numeral 330. The operations of the process 400 are described generally as being performed by the system 200. The operations of the process 400 may be performed exclusively by the data processing system 205, may be performed exclusively by the database system 280, or may be performed by a combination of the data processing system 205 and the database system 280. In some implementations, operations of the process 400 may be performed by one or more processors included in one or more electronic devices.

The system 200 determines a number of distinct data values within each of at least two attribute classes (410). For example, the system 200 may process data within each of at least two attribute classes by analyzing each data value for an attribute class and counting the number of distinct values that exist for the attribute class in the data. In this example, the system 200 may sequentially analyze all of the records in the data, track data values for the attribute class present in the data

records (e.g., store analyzed values in temporary storage), and compare data values for subsequent records to the tracked values to determine whether the data values are distinct from other data values included in the data records. When a data value matches a tracked value, the system 200 determines that the data value is not distinct (e.g., determines that the data value is redundant of at least one other data value) and continues processing the next data record without updating tracked data. When a data value does not match any tracked value, the system 200 determines that the data value is distinct (e.g., determines that the data value is not redundant of at least one other data value), stores the data value with the tracked data values for comparison against subsequent records, and increments a counter that tracks the number of distinct data values within the attribute class.

In some implementations, the system 200 may sort the data with respect to an attribute class of interest prior to identifying the number of distinct values within the attribute class of interest. Sorting the accessed data may improve efficiency in identifying the number of distinct values because the data records with the same data value for the attribute class would be arranged together and processed consecutively. Accordingly, because the system 200 knows the data values are arranged consecutively, the system 200 may only have to compare a data value to the most recently tracked data value. Specifically, if the data value is redundant of a previously processed data value, it is necessarily redundant of the most recently tracked data value because it would have been grouped together with the most recently tracked data value in the sorting process.

The system 200 determines a parameter for at least one of the distinct data values that reflects contribution of the distinct data value to the entirety of the data values in an attribute class (420). For a particular distinct data value, the system 200 may determine a parameter that indicates the number of times that the particular distinct data value is found within the data records or may determine a parameter that indicates the percentage of the data records in which the particular distinct data value is found. The system 200 also may determine a rate of occurrence of the distinct data value within the entirety of data values within the attribute class. The system 200 may determine a parameter that reflects contribution of the distinct data value to the entirety of the data values in an attribute class for each of the distinct data values in the attribute classes (e.g., each distinct value is associated with a parameter). Using the parameters, the system 200 may determine a distribution of redundant data within an attribute class.

The system 200 determines a redundancy measure for the attribute class based on the number of distinct data values and the determined parameter (430). The system 200 may apply the number of distinct data values and the determined parameter to a formula that computes the redundancy measure. For instance, the system 200 may use the determined parameter as a weighting value in evaluating the number of distinct data values. The system 200 may apply a weighting value that increases a measured level of data redundancy when the parameter reflects a relatively high contribution of a redundant data value to the entirety of the data values in an attribute class. The system 200 also may apply a weighting value that decreases a measured level of data redundancy when the parameter reflects a relatively low contribution of a redundant data value to the entirety of the data values in an attribute class. In some examples, the system 200 may determine a redundancy measure that reflects a relatively higher level of data redundancy when the parameter reflects a relatively higher contribution of the distinct data value to the entirety of the data values in the attribute class. In these examples, the

11

system 200 may determine a redundancy measure that reflects a relatively lower level of data redundancy when the parameter reflects a relatively lower contribution of the distinct data value to the entirety of the data values in the attribute class. By using the parameter, the system 200 may account for the ability of the system 200 to leverage redundancy within the data, rather than just the number of distinct data values.

The system 200 determines an order for organizing the attribute classes of the data based on the redundancy measure (440). For example, the system 500 may compare determined redundancy measures for each of the attribute classes and determine an order in which to organize the attribute classes based on the comparison. In this example, the system 200 may select the attribute class with a redundancy measure that reflects an ability to leverage data redundancy to a highest degree as the first attribute class in the order. The system 200 may order the remaining attribute classes in an order of redundancy measures that reflect a decreasing ability to leverage data redundancy, perhaps computing new redundancy measures after selecting an attribute class in the order.

In some examples, the system 200 may determine that a first attribute class has a first number of distinct data values and a first parameter for a first distinct data value. The first parameter for the first distinct data value may reflect a relatively high contribution of the first distinct data value to the entirety of data values within the first attribute class (e.g., the parameter may reflect a ninety percent contribution of the first distinct data value). The system 200 also may determine that a second attribute class has a second number of distinct data values and a second parameter for a second distinct data value. In these examples, the second number of distinct data values within the second attribute class is lower than the first number of distinct data values within the first attribute class and the second parameter for the second distinct data value may reflect a relatively low contribution of the second distinct data value to the entirety of data values within the second attribute class (e.g., the parameter may reflect a five percent contribution of the second distinct data value). The system 200 may determine a first redundancy measure for the first attribute class and a second redundancy measure for the second attribute class. Because the first parameter reflects a relatively high contribution of the first distinct data value to the entirety of data values within the first attribute class and the second parameter for the second distinct data value reflects a relatively low contribution of the second distinct data value to the entirety of data values within the second attribute class, the second redundancy measure may reflect a lower level of redundancy for the second attribute class than the first redundancy measure reflects for the first attribute class despite the second number of distinct data values within the second attribute class being lower than the first number of distinct data values within the first attribute class. The system 200 may order the first attribute class prior to the second attribute class in the determined relative order based on the second redundancy measure reflecting a lower level of redundancy for the second attribute class than the first redundancy measure reflects for the first attribute class. Ordering the first attribute class prior to the second attribute class may enable the system 200 to leverage redundancy of data better than if the second attribute class was ordered prior to the first attribute class because of the relatively high contribution of the first distinct data value to the first attribute class, even though the second attribute class has fewer distinct data values.

FIG. 5 illustrates a process 500 for determining an order for organizing multiple attribute classes of data based on identified redundancy characteristics. The process 500 may be used in

12

determining an order for organizing multiple attribute classes of data based on identified redundancy characteristics referenced above with respect to reference numeral 330. The operations of the process 500 are described generally as being performed by the system 200. The operations of the process 500 may be performed exclusively by the data processing system 205, may be performed exclusively by the database system 280, or may be performed by a combination of the data processing system 205 and the database system 280. In some implementations, operations of the process 500 may be performed by one or more processors included in one or more electronic devices.

The system 200 accesses one or more rules that define secondary criteria related to organizing attribute classes of data (510). For example, the system 200 may access, from electronic storage, rules that define organization of attribute classes of data. The rules may define relationships of attribute classes that are to be maintained when determining an order for the attribute classes and may define preferences, which are unrelated to redundancy of data, that impact ordering of attribute classes. The rules also may define thresholds related to data redundancy that are used in determining whether to reorder or reorganize attribute classes within a set of data. The rules may be general rules that are applied generally to any operation to determine an order for organizing multiple attribute classes of data or may be specific to a specific set of data or one or more specific attribute classes. The rules may be routinely used in determining an order for attribute classes or may be used in determining an order for attribute classes to break a tie in redundancy measures or when redundancy analysis of the data within the attribute classes is inconclusive.

In some implementations, the system 200 may access a rule that indicates a preference to maintain an order of attribute classes within a dimension of attribute classes despite redundancy characteristics. A dimension of attribute classes may define a subset of multiple attribute classes that have a parent-child relationship (e.g., a dimension of time may have a parent class of "year" and child classes of "month," "day," and "hour"). The rule may prevent attribute classes within a dimension from being reordered. For example, in the dimension of time, the rule may prevent the system 200 from ordering the attribute class "day" prior to the attribute class "year," even when redundancy characteristics of data values in the "year" and "day" attribute classes indicate that the "day" attribute class has a higher degree of data redundancy than the "year" attribute class.

The rule also may allow reordering of attribute classes within a dimension, but require a relatively high difference in data redundancy to do so. For instance, in the dimension of time, the rule may require the system 200 to determine that redundancy measures for the "year" and "day" attribute classes have at least a first threshold difference to enable the system 200 to order the attribute class "day" prior to the attribute class "year." The first threshold difference may be greater than a second threshold difference used by the system 200 in determining whether to reorder attribute classes that are not part of the same dimension. In some examples, the system 200 may access a rule that defines multiple thresholds for determining whether to reorder attribute classes within a dimension, where the multiple thresholds are based on relative position of attribute classes within the dimension. In these examples, threshold differences may increase as the relative distance between attribute classes within the dimension increases. For the time dimension, a first threshold difference may be required to reorder the attribute class "month" prior to the attribute class "year" and a second threshold

13

difference that is greater than the first threshold difference may be required to reorder the attribute class “hour” prior to the attribute class “year.”

In some implementations, the system **200** may access a rule that indicates other characteristics to consider in ordering attribute classes. For instance, the system **200** may access a rule that indicates that a relatively frequently searched attribute class is prioritized in comparison to a relatively infrequently searched attribute class in determining an order of attribute classes. The rule may prevent the system **200** from ordering a relatively frequently searched attribute class prior to a relatively infrequently searched attribute class when the difference in the amount with which the relatively frequently searched attribute class and the relatively infrequently searched attribute class exceeds a threshold. In this regard, when a particular attribute class is often searched, the system **200** may ensure that redundancy of the particular attribute class is leveraged. Although redundancy of another attribute class that is infrequently searched may be leveraged to a greater degree (e.g., a higher level of compression), because the other attribute class is not searched often, the benefit of leveraging the redundancy is relatively low as compared to the benefit leveraging the redundancy of the particular attribute class. Specifically, the aggregated benefit of leveraging the redundancy of the particular attribute class over many searches may be greater than the aggregated benefit of leveraging the redundancy of the other attribute class over few searches, even though the benefit of a single instance is greater for the other attribute class.

The system **200** also may access a rule that defines a threshold difference in redundancy measures to use in ordering attribute classes based on other characteristics. For example, the system **200** may determine a threshold difference that varies based on how frequently two attribute classes are searched. In this example, when a relatively frequently searched attribute class is being compared to a relatively infrequently searched attribute class, the rule may require a first threshold difference in redundancy measures in order to order the relatively infrequently searched attribute class prior to the relatively frequently searched attribute class. When a first relatively frequently searched attribute class is being compared to a second relatively frequently searched attribute class, the rule may require a second threshold difference in redundancy measures that is less than the first threshold difference in order to order the second relatively frequently searched attribute class prior to the first relatively frequently searched attribute class.

The system **200** analyzes secondary criteria of the attribute classes of the data based on the one or more rules (**520**). For instance, the system **200** may identify second criteria of the attribute classes being considered and compare the identified secondary criteria to the one or more rules. Based on the comparison, the system **200** may evaluate the secondary criteria and update a weighting value or measure that is used in ordering the attribute classes. The system **200** also may store data that is later used in ordering the attribute classes and that defines relationships between the attribute classes based on the evaluation of the secondary criteria.

When the rule is related to handling a dimension of attribute classes, the system **200** may identify a dimension of attribute classes included in the multiple attribute classes. The system **200** may identify one or more dimensions of attribute classes based on the attribute classes being arranged in a particular order in the electronic data storage based on a parent-child relationship. The system **200** may identify one or more dimensions of attribute classes within the attribute classes under consideration (e.g., which attribute classes are

14

in which dimensions) and also may identify relative positions of attribute classes within each of the identified dimensions. Based on the identified dimensions and the identified relative positions of attribute classes within each of the identified dimensions, the system **200** may apply the rule related to handling a dimension of attribute classes to the attribute classes under consideration.

When the rule is related to other characteristics to consider in ordering attribute classes, the system **200** may identify other characteristics for the multiple attribute classes. For example, when the other characteristics relate to how frequently the attribute classes are searched, the system **200** may access data related to how frequently the attribute classes are searched. In this example, the system **200** may track searches over time and store data that reflects how frequently each of the attribute classes are searched. The system **200** may store data tracking the number of times an attribute class is searched for a given period of time and compute a frequency with which the attribute class is searched by dividing the number of times by the given period of time. Using the frequency of how often each of the attribute classes is searched, the system **200** may identify relatively frequently searched attribute classes and relatively infrequently searched attribute classes and evaluate the frequency with which the attribute classes are searched with respect to the rule.

The system **200** determines a redundancy measure for each of the attribute classes (**530**). The system **200** may determine a redundancy measure for each of the attribute classes using any of the techniques described through the disclosure for evaluating redundancy of data within an attribute class.

The system **200** determines an order for organizing the attribute classes of the data based on the analysis of the secondary criteria and the redundancy measures (**540**). The system **200** may determine an order that accounts for the redundancy measures, but also complies with the one or more rules. For instance, the system **200** may set thresholds to use in evaluating redundancy measures based on the analysis of the one or more rules and then determine an order for the multiple attribute classes by evaluating the redundancy measures with respect to the set thresholds.

When the rule is related to handling a dimension of attribute classes, the system **200** may maintain the particular order of the attribute classes included in the dimension despite the identified redundancy characteristics indicating that a different order of the attribute classes included in the dimension is preferred. For instance, the rule may prevent the system **200** from reordering attribute classes within a dimension and the system **200** may maintain the particular order of the attribute classes included in the dimension to comply with the rule. In another example, the rule may allow the system **200** to reorder attribute classes within a dimension, but require a higher threshold difference in redundancy to do so. In this example, the system **200** may compute a difference in redundancy measures between a first attribute class and a second attribute class within a dimension, determine that the difference in redundancy measures is less than the threshold difference, and order the first attribute class prior to the second attribute class, even though the redundancy measure of the second attribute class indicates that the second attribute class has a higher degree of redundancy than the first attribute class.

When the rule is related to how frequently the attribute classes are searched, the system **200** may order the attribute classes based on how frequently the attribute classes are searched. For instance, based on an identification of a first attribute class as a relatively frequently searched attribute class and a second attribute class as a relatively infrequently

15

searched attribute class, the system **200** may determine to order the first attribute class prior to the second attribute class despite the identified redundancy characteristics indicating that ordering the second attribute class prior to the first attribute class is preferred. The rule may prevent the system **200** from ordering a first attribute class prior to a second attribute class when a difference in search frequencies of the first and second attribute classes exceeds a threshold. The system **500** may compute a difference between a first search frequency for a first attribute class and a second search frequency for a second attribute class, determine that the difference in search frequencies exceeds the threshold, and order the first attribute class prior to the second attribute class, even though the redundancy measure of the second attribute class indicates that the second attribute class has a higher degree of redundancy than the first attribute class.

In another example, the rule may require a higher threshold difference in redundancy to order a relatively infrequently search attribute class prior to a relatively frequently search attribute class. In this example, the system **200** may compute a difference in redundancy measures between a first attribute class that is relatively frequently searched and a second attribute class that is relatively infrequently searched, determine that the difference in redundancy measures is less than the threshold difference, and order the first attribute class prior to the second attribute class, even though the redundancy measure of the second attribute class indicates that the second attribute class has a higher degree of redundancy than the first attribute class.

Referring again to FIG. 3, the system **200** organizes the attribute classes of the accessed data based on the determined order (**340**). For instance, the system **200** arranges the data in a manner that reflects the determined order of the attribute classes. The system **200** may rearrange the data accessed from the electronic data storage and thereby reorder the attribute classes. When the data accessed from the electronic data storage is part of a database table with rows and columns, the system **200** may reorder the columns within the database table. In organizing the attribute classes of the accessed data based on the determined order, the system **200** may physically rearrange the accessed data into different storage locations based on the determined order. The system **200** also may leave the accessed data in the same storage locations, but update a data structure that logically represents an order of the accessed data to reflect the determined order. For example, the system **200** may use a database table index that defines an order of columns of data within the database table and maps the columns of data to physical storage locations. In this example, the system **200** may update the database table index to reflect a new order for the columns in the database table while maintaining the mappings of the columns of data to physical storage locations.

In some examples, the system **200** may modify or change the order of the attribute classes in which the data is organized in the electronic data storage. When the multiple attribute classes of the accessed data are organized in the electronic data storage in a first order, the system **200** may reorganize the multiple attribute classes of the accessed data in a second order that is different than the first order in which the multiple attribute classes of the accessed data are organized in the electronic data storage.

The system **200** also may reorganize the attribute classes in a manner that is different from the dimensional relationships by which the attribute classes are organized in the electronic data storage. For example, the multiple attribute classes may include at least a first dimension of attribute classes and a second dimension of attribute classes. In this example, each

16

of the first and second dimensions define an exclusive subset of the multiple attribute classes that are related. The multiple attribute classes of the accessed data further may be organized in the electronic data storage in a first order that is based on the first and second dimensions such that the exclusive subset of attribute classes included in the first dimension are ordered consecutively and the exclusive subset of attribute classes included in the second dimension are ordered consecutively. In this example, the system **200** may reorganize the multiple attribute classes of the accessed data in a second order in which at least one attribute class included in the first dimension is ordered among the attribute classes included in the second dimension such that the exclusive subset of attribute classes included in the first dimension are no longer ordered consecutively and the exclusive subset of attribute classes included in the second dimension are no longer ordered consecutively.

To illustrate, suppose the accessed data includes a time dimension with the attribute classes "year," "month," "day," and "hour," and also includes a location dimension that includes the attribute classes "state," "city," and "street address." The accessed data may be stored in a database table with a column corresponding to each of the attribute classes. The columns may be arranged in an order "year," "month," "day," "hour," "state," "city," and "street address." After determining an order of the attribute classes based on identified redundancy characteristics, the system **200** may determine that the shown order does not best leverage redundancy within the data and a different order that disrupts the dimensional arrangement is preferred. For instance, the system **200** may determine that an order of "year," "state," "city," "month," "day," "hour," and "street address" best leverages redundancy within the accessed data and reorder the attribute classes in this order. In this case, the system **200** has interspersed attribute classes from a first dimension within attribute classes of a second dimension, but maintained the relative order of attribute classes within each of the respective dimensions (i.e., if the attribute classes from the first dimension are removed, the attribute classes from the second dimension remain in the same order as prior to reordering the attribute classes).

In some examples, the system **200** may reorder the attribute classes without regard for dimension relationships. In this regard, the system **200** may reorder attribute classes within a dimension in an order that is different than the order defined by the dimension. For instance, the multiple attribute classes may include a dimension of attribute classes that defines a subset of the multiple attribute classes that have a parent-child relationship and the multiple attribute classes of the accessed data may be organized in the electronic data storage in a first order that is based on the dimension such that parent attribute classes are ordered prior to child attribute classes in the electronic data storage. The system **200** may reorganize the subset of the multiple attribute classes included in the dimension in a second order in which at least one child attribute class is ordered prior to at least one of its parent attribute classes as defined by the parent-child relationship.

To illustrate, suppose the accessed data includes a time dimension with the attribute classes "year," "month," "day," and "hour." The accessed data may be stored in a database table with a column corresponding to each of the attribute classes. The columns may be arranged in an order "year," "month," "day," and "hour." After determining an order of the attribute classes based on identified redundancy characteristics, the system **200** may determine that the shown order does not best leverage redundancy within the data and a different order that disrupts an order of the time dimension is preferred.

17

For instance, the system 200 may determine that an order of “month,” “year,” “day,” and “hour” best leverages redundancy within the accessed data and reorder the attribute classes in this order, even though the order defined by the time dimension is changed. This type of change may occur when the data spans a relatively large number of years, but with events occurring in a relatively low number of months. For example, tax return filing data may include filings that span many years (e.g., a period of 20 years), but with the filings typically occurring in a low number of months or occurring with a relatively high frequency in a single month (e.g., April).

The system 200 compresses the data within the organized attribute classes using run length encoding (350). Run length encoding is a form of lossless data compression in which runs of data (e.g., sequences in which the same data value occurs in many consecutive data elements) are stored as a single data value and count, rather than as a series of repeated values. Run length encoding may provide a significant level of compression when the data has a large amount of redundancy and a data value for many records may be expressed as a single value and range, instead of storing the value for each record. In other implementations, other forms of compression (or no compression) may be used.

In some implementations, after the attribute classes of the accessed data have been organized, the system 200 may sort the data to group the same values within the first attribute class together. The system 200 then may store each group of the same value as a single block of data within the first attribute class. Each block may include data that defines the value associated with the block and data that defines the records (e.g., rows) associated with the block. Because the block does not redundantly store the value for each record (e.g., row), the data may be compressed and the storage capacity needed to store the data may be reduced.

After storing the data within the first attribute class as compressed blocks, the system 200 may compress data values within each block for the second attribute class in the determined order. The system 200 may sort the data values to group the same value for the second attribute class within a block together and then store the grouped values as compressed blocks. Because the first attribute class has been compressed into blocks when the system 200 evaluates the second attribute class, the structure of the data with respect to the first attribute class has been set and, therefore, the system 200 cannot reorder or rearrange records (e.g., rows) within a first block of the first attribute class to a second block of the first attribute class. Accordingly, even though two records (e.g., rows) may have a redundant value for the second attribute class, the system 200 may not be able to take advantage of that redundancy and compress the two values into a single block of common values, unless the records (e.g., rows) are arranged in the same block in the first attribute class. The system 200 may continue to process blocks within each attribute class until all of the attribute classes have been processed or until an attribute class does not include any compressed blocks of data that span multiple records (e.g., rows).

In some implementations, the system 200 may load a portion of the data from the electronic data storage and process the portion of the data as described above. After compressing the portion of the data, the system 200 may determine how much storage space is needed to store the compressed data and how much storage space the system 200 has in which to store data (e.g., a size of a memory). The system 200 may compare the two values and determine whether additional data may be accessed and loaded into memory. For instance,

18

when the data is highly redundant and a large amount of compression occurs, the system 200 may determine that a relatively large amount of storage space remains. When a relatively large amount of storage space remains, the system 200 may load more data from the electronic data storage (e.g., database) to take advantage of the additional storage space gained through data compression.

FIG. 6 illustrates an example data structure as the data within the data structure is organized by attribute classes and compressed using run length encoding. As shown, the data structure 610 illustrates how the data in this example is stored in the electronic data storage. The data structure 610 includes a month column 611 that stores a month value associated with a record, a year column 612 that stores a year value associated with a record, a state column 613 that stores a state value associated with a record, a city column 614 that stores a city value associated with a record, and a sales metric column 615 that stores a sales metric associated with a record. The data structure 610 stores the data with the attribute classes in the order 611, 612, 613, 614, and 615. The data structure 610 also includes data records 621 to 627 that each include data associated with a stored record.

The system 200 first processes the data structure 610 to identify redundancy characteristics of the data within the data structure 610. For example, the system 200 may count the number of distinct data values in each attribute class. In this example, the system 200 determines that the year column 612 has two distinct values, the month column 611 has three distinct values, the state column 613 has three distinct values, the city column 614 has four distinct values, and the sales metric column 615 has seven distinct values. Because the year column 612 has the fewest number of distinct data values, the system 200 determines to order the year column 612 first.

In response to the determination to order the year column 612 first, the system 200 identifies redundancy characteristics of the other columns with respect to blocks of data within the year column 612. Because the year column 612 has two distinct data values, two blocks of data exist: the “06” block and the “07” block. In identifying redundancy characteristics of the other columns within the blocks of the year column 612, the system 200 determines that the month column 611 has two distinct values in the “06” block and one distinct data value in the “07” block for a total of three distinct data values, the state column 613 has two distinct values in the “06” block and three distinct data values in the “07” block for a total of five distinct data values, the city column 614 has three distinct values in the “06” block and three distinct data values in the “07” block for a total of six distinct data values, and the sales metric column 615 has four distinct values in the “06” block and three distinct data values in the “07” block for a total of seven distinct data values. Because the month column 611 has the fewest number of distinct data values with respect to blocks of data within the year column 612, the system 200 determines to order the month column 611 second in the attribute classes.

In response to the determination to order the month column 611 second, the system 200 identifies redundancy characteristics of the remaining columns with respect to blocks of data within the month column 611. Because the month column 611 has three distinct data values with respect to blocks of data within the year column 612, three blocks of data exist: the “06/01” block, the “06/02” block, and the “07/01” block. In identifying redundancy characteristics of the remaining columns within the blocks of the month column 611, the system 200 determines that the state column 613 has one distinct value in the “06/01” block, two distinct values in the “06/02” block and three distinct values in the “07/01” block

19

for a total of six distinct data values, the city column **614** has one distinct value in the “06/01” block, three distinct values in the “06/02” block and three distinct values in the “07/01” block for a total of seven distinct data values, and the sales metric column **615** has one distinct value in the “06/01” block, three distinct values in the “06/02” block and three distinct values in the “07/01” block for a total of seven distinct data values. Because the state column **612** has the fewest number of distinct data values with respect to blocks of data within the month column **611**, the system **200** determines to order the state column **612** next (e.g., third) in the attribute classes.

In addition, because the city column **614** and the sales metric column **615** both have seven distinct data values with respect to blocks of data within the month column **611** and only seven records exist in the accessed data, the system **200** determines that no additional data redundancy exists in those columns. Based on the determination that no additional data redundancy exists in the remaining columns, the system **200** maintains the order of the remaining columns in the determined order. In some implementations, the system **200** may have determined that the sales metric column **615** had no data redundancy (e.g., seven distinct data values) when identifying the first column and, therefore, set the sales metric column **615** last in the determined order and stopped processing the sales metric column **615** in determining the order. In other implementations, because the sales metric column **615** includes sales metric data that is the subject of a report being generated using the accessed data, the sales metric column **615** may not be considered in ordering the attribute classes and only considered when computing the relevant values in generating a report.

Accordingly, the system **200** determines the final order for the attribute classes as columns **612**, **611**, **613**, **614**, and **615** based on the identified redundancy characteristics. Using the determined order, the system **200** reorganizes the attribute classes in the data structure **610** as shown in the data structure **630**. In addition, the system **200** has sorted the data in the data structure **630** into groups of data records having the same value. After the first column (i.e., the year column **612**), the system **200** sorted the data within blocks included in the immediately preceding column in the determined order.

After the data has been organized and sorted, the system **200** compresses the data into blocks using run length encoding. The data structure **640** illustrates the results of the compression using run length encoding. As shown, the data structure **640** includes a block **641** that represents the “06” data value in the year column **612** and a block **642** that represents the “07” data value in the year column **612**. The data structure **640** also includes a block **643** that represents the “06/01” data value in the month column **611** for the block **641**, a block **644** that represents the “06/02” data value in the month column **611** for the block **641**, and a block **645** that represents the “07/01” data value in the month column **611** for the block **642**. The data structure **640** further includes a block **646** that represents the “VA” data value in the state column **613** for the block **643**. As such, the data structure **640** represents an example of data after the data has been reorganized and compressed into blocks using run length encoding.

Referring again to FIG. 3, the system **200** generates an index that is descriptive of the compressed data (**360**). The system **200** may analyze the compressed data and generate a data structure that defines the organization of the data. For instance, when the data is organized and compressed into blocks, the index may identify blocks within the organized data and identify relationships between the blocks in the organized data. The relationships may indicate whether

20

blocks are within the same attribute class and whether a particular block is related to other blocks within other attribute classes (e.g., whether the block has a parent block associated with a parent attribute class and/or whether the block has a child block associated with a child attribute class). The system **200** may generate the index by identifying the blocks within the organized and compressed data, determining relationships between the identified blocks, and generating data that is descriptive of the identified blocks and determined relationships. The system **200** also may identify storage locations of data records associated with the identified blocks and store addressing information (e.g., row numbers) for each of the identified blocks that map the identified blocks to stored records. The addressing information may enable future accesses of data stored within records associated with the identified blocks by referencing the index.

In some implementations, the system **200** may generate the index by identifying blocks within the compressed data that have common values and identifying storage locations of the identified blocks within the compressed data. The system **200** may associate, within the index, the identified blocks within the compressed data with the corresponding common values and the identified storage locations. By associating an identified block within the compressed data with the corresponding common value and the identified storage locations, the system **200** may be able to group data having the common value and identify storage locations associated with records having the common value relatively quickly using the index.

FIGS. 7 and 8 illustrate example data structures of an index that is descriptive of the compressed data shown in FIG. 6. FIG. 7 illustrates a graph structure **700** in which each block in the compressed data is represented by a node and the arrows between the nodes represent relationships between the blocks within the compressed data. As shown, the node **710** corresponds to the “06” block in the year column **612** (e.g., the block **641**). The node **710** is connected to the nodes **720** and **730** because the nodes **720** and **730** are immediate child nodes or blocks associated with the node **710**. The node **720** corresponds to the “06/01” block (e.g., the block **643**) in the month column **611** for the “06” block in the year column **612**, and the node **730** corresponds to the “06/02” block (e.g., the block **644**) in the month column **611** for the “06” block in the year column **612**. The node **720** has one immediate child node **721** for the “VA” data value in the state column **613** and the node **721** has one immediate child node **722** for the “McLean” data value in the city column **614**. The node **730** has a first immediate child node **731** for the “NY” data value in the state column **613** and a second immediate child node **732** for the “VA” data value in the state column **613** (e.g., the block **646**). The node **731** has one immediate child node **733** for the “New York” data value in the city column **614**. The node **732** has a first immediate child node **734** for the “McLean” data value in the city column **614** and a second immediate child node **735** for the “Vienna” data value in the city column **614**.

The node **740** corresponds to the “07” block in the year column **612** (e.g., the block **642**) and shows relationships of the blocks in the data associated with records having “07” as the year value. The node **740** is connected to the node **750** because the node **750** is an immediate child node or block associated with the node **740**. The node **750** corresponds to the “07/01” block (e.g., the block **645**) in the month column **611** for the “07” block in the year column **612**. The node **750** has a first immediate child node **751** for the “NY” data value in the state column **613**, a second immediate child node **752** for the “MD” data value in the state column **613**, and a third immediate child node **753** for the “VA” data value in the state column **613**. The node **751** has one immediate child node **754**



21

for the “New York” data value in the city column **614**, the node **752** has one immediate child node **755** for the “Baltimore” data value in the city column **614**, and the node **753** has one immediate child node **756** for the “Vienna” data value in the city column **614**.

FIG. **8** illustrates an element index **800**. The element index **800** includes a column for each of the attribute classes in the accessed data. For instance, the element index **800** includes a year column **810**, a month column **820**, a state column **830**, and a city column **840**. The attribute classes in the element index **800** are arranged in the order of the attribute classes after the attribute classes have been ordered based on identified redundancy characteristics. For each attribute class, the element index **800** includes one entry for each distinct data value within the corresponding attribute class. For each distinct data value, the element index **800** includes the distinct data value and block identification data that identifies which one or more blocks within the attribute class have the corresponding data value.

The element index **800** also is associated with tables that correspond to attribute classes included in the accessed data. The tables store a block identifier for each block included in an attribute class and store location information for data records that are included in the corresponding block. The location information enables the system **200** to identify and access data records that are included in a block relatively quickly. As shown, the table **850** corresponds to the year attribute class. The table **850** includes an entry **851** corresponding to the first block (e.g., Block **0**) in the year attribute class and entry **852** corresponding to the second block (e.g., Block **1**) in the year attribute class. The entries **851-852** define which rows in a database table correspond to the block identified in the corresponding entry and the table **850** includes data associating each row in the database with a block.

The table **860** corresponds to the month attribute class. The table **860** includes an entry **861** corresponding to the first block (e.g., Block **0**) in the month attribute class, an entry **862** corresponding to the second block (e.g., Block **1**) in the month attribute class, and an entry **863** corresponding to the third block (e.g., Block **2**) in the month attribute class. The entries **861-863** define which rows in a database table correspond to the block identified in the corresponding entry and the table **860** includes data associating each row in the database with a block.

The table **870** corresponds to the state attribute class. The table **870** includes an entry **871** corresponding to the first block (e.g., Block **0**) in the state attribute class, an entry **872** corresponding to the second block (e.g., Block **1**) in the state attribute class, an entry **873** corresponding to the third block (e.g., Block **2**) in the state attribute class, an entry **874** corresponding to the fourth block (e.g., Block **3**) in the state attribute class, an entry **875** corresponding to the fifth block (e.g., Block **4**) in the state attribute class, and an entry **876** corresponding to the sixth block (e.g., Block **5**) in the state attribute class. The entries **871-876** define which rows in a database table correspond to the block identified in the corresponding entry and the table **870** includes data associating each row in the database with a block.

The table **880** corresponds to the city attribute class. The table **880** includes an entry **881** corresponding to the first block (e.g., Block **0**) in the city attribute class, an entry **882** corresponding to the second block (e.g., Block **1**) in the city attribute class, an entry **883** corresponding to the third block (e.g., Block **2**) in the city attribute class, an entry **884** corresponding to the fourth block (e.g., Block **3**) in the city attribute class, an entry **885** corresponding to the fifth block (e.g., Block **4**) in the city attribute class, an entry **886** corre-

22

sponding to the sixth block (e.g., Block **5**) in the city attribute class, and an entry **887** corresponding to the seventh block (e.g., Block **6**) in the city attribute class. The entries **881-887** define which rows in a database table correspond to the block identified in the corresponding entry and the table **880** includes data associating each row in the database with a block.

The system **200** may use the element index **800** to identify rows in the database that have a particular data value for a particular attribute class. For instance, the system **200** may compare the particular data value for the particular attribute class to the element index **800** to identify with which one or more blocks the particular data value is associated. Once the system **200** has identified the one or more blocks associated with the particular data value, the system **200** may access the table corresponding to the particular attribute class, use the identified one or more blocks as identifiers to locate relevant entries within the accessed table, and access the storage location information (e.g., row numbers) included in the located entries. The system **200** may use the storage location information to access data records that have the particular data value for the particular attribute class. Because the system **200** uses the linked addressing in the element index **800** to identify the relevant data records, the speed of identifying and accessing the relevant data records may be faster than a system that directly searches the database table to identify records that have the particular data value for the particular attribute class.

To illustrate, suppose the system **200** is searching for data records that have a value “NY” for the state attribute class. The system **200** accesses the element index **800**, compares the value “NY” to values in the state column **830**, and, based on the comparison, determines that the value “NY” corresponds to blocks **1** and **3** in the state attribute class. Based on the identification of blocks **1** and **3**, the system **200** accesses the table **870** and identifies the storage location information correspond to blocks **1** and **3** as rows **1** and **4**. Accordingly, the system **200** identified all of the rows in the database table that correspond to the value “NY” in the state attribute class without having to directly search records in the database table.

Referring again to FIG. **3**, the system **200** stores, in electronic storage, the compressed data and the generated index to enable subsequent searching of the compressed data using the generated index (**370**). For instance, the system **200** may store the compressed data and the generated index in the memory **260**. By storing the compressed data and the generated index in the memory **260**, which the data processing system **205** is able to access more quickly than the database system **280**, the system **200** may generate reports and execute queries in a relatively fast manner by accessing the index from the memory **260** and using the index **260** to identify relevant portions of the compressed data.

In some implementations, the system **200** may handle a request to access a particular block within the compressed data using the index. For example, the system **200** may access, from the electronic storage, the generated index and identify the particular block within the generated index. In this example, the system **200** may identify, using the index, particular storage locations corresponding to the particular block and access, from the electronic storage, data (e.g., metric values) corresponding to the particular storage locations identified using the index.

The system **200** may use the compressed data and the generated index to generate a report. For example, the system **200** may receive a report generation query that defines a subset of the multiple attribute classes of interest. In response



23

to receiving the report generation query, the system 200 may access, from electronic storage (e.g., the memory 260), the generated index and identify the subset of the multiple attribute classes of interest defined by the report generation query. Using the generated index, the system 200 may identify portions of the compressed data that include a distinct combination of values for the subset of the multiple attribute classes of interest and access, from electronic storage (e.g., the memory 260), metrics (e.g., sales data) for each of the identified portions of the compressed data. Based on the accessed metrics, the system 200 may compute a report parameter (e.g., average sales data) for each of the identified portions of the compressed data and generate a report based on the computed report parameters. The report may list each distinct combination of values for the subset of the multiple attribute classes of interest and the computed report parameter associated with the corresponding distinct combination. The system 200 may display, on a display device, the generated report responsive to the report generation query.

FIG. 9 illustrates a process 900 for encoding data to enable subsequent searching of the encoded data. The operations of the process 900 are described generally as being performed by the system 200. The operations of the process 900 may be performed exclusively by the data processing system 205, may be performed exclusively by the database system 280, or may be performed by a combination of the data processing system 205 and the database system 280. In some implementations, operations of the process 900 may be performed by one or more processors included in one or more electronic devices.

The system 200 accesses, from an electronic data storage, a data set that includes information for multiple attribute classes (910). The system 200 may access the data set that includes information for multiple attribute classes using techniques similar to those described above with respect to reference numeral 310 in FIG. 3.

The system 200 identifies redundancy characteristics of the data set within each of at least two attribute classes (920). The system 200 may identify redundancy characteristics of the data set within each of at least two attribute classes using techniques similar to those described above with respect to reference numeral 320 in FIG. 3.

The system 200 determines an order for organizing the multiple attribute classes of the data set based on the identified redundancy characteristics (930). The system 200 may determine an order for organizing the multiple attribute classes of the data set using techniques similar to those described above with respect to reference numeral 330 in FIG. 3. The system 200 orders the attribute classes of the data set based on the determined order (940). The system 200 may order the attribute classes of the data set based on the determined order using techniques similar to those described above with respect to reference numeral 340 in FIG. 3.

The system 200 encodes the data within the attribute classes, with the encoding reflecting redundancies and uniqueness within the data set and also accounting for the order applied to the attribute classes (950). For example, the system 200 may identify the first attribute class in the order, sort the data set based on values in the first attribute class, and encode the values in the first attribute class using run length encoding. In this example, the system 200 may identify the next attribute class in the order and identify portions of the data set (e.g., sets of data records included in the data set) that correspond to the groups of identical data values resulting from the run length encoding performed on the first attribute class. For each identified portion of the data set, the system 200 may independently sort and run length encode the values

24

in the next attribute class. The system 200 may continue to encode the data in the remaining attribute classes using the process used to encode the data in the next attribute class. In some examples, the system 200 may encode the data within the attribute classes using techniques similar to those described above with respect to reference numeral 350 in FIG. 3.

The system 200 stores, in electronic storage, the encoded data to enable subsequent searching of the encoded data (960). The system 200 may store the encoded data to enable subsequent searching of the encoded data using techniques similar to those described above with respect to reference numeral 370 in FIG. 3.

FIG. 10 illustrates a process 1000 for generating a report based on a query. The operations of the process 1000 are described generally as being performed by the system 200. The operations of the process 1000 may be performed exclusively by the data processing system 205, may be performed exclusively by the database system 280, or may be performed by a combination of the data processing system 205 and the database system 280. In some implementations, operations of the process 1000 may be performed by one or more processors included in one or more electronic devices.

The system 200 accesses an index of data based on receipt of a query entered in an input interface 1010 (1020). For instance, the system 200 may display an input interface 1010 to a user. The user may enter a query into the input interface 1010 and the system 200 may receive and process the inputted query. The query may define a computational metric in which the user is interested and the attribute classes the user wishes to use to group the data and compute the metric. As shown, the query entered in the input interface 1010 defines the computational metric of average sales metric and the relevant attribute classes as year and state. Accordingly, in response to receipt of the query shown in the input interface 1010, the system 200 groups the data into groups of records with common values for the year attribute and the state attribute and computes an average sales value for each group.

To group the data and compute the metric, the system 200 accesses an index that is descriptive of the data that is the subject of the user's query. For example, the system 200 may access an index from electronic storage (e.g., memory) that is descriptive of data that has been organized and compressed as discussed above with respect to FIGS. 1-9. In this example, the system 200 may access the index shown in FIGS. 7 and 8. The system 200 may automatically identify the appropriate index to access in response to the query. The system 200 may only maintain a single index and, therefore, the system 200 may access the same index for any received query.

When the system 200 stores multiple indexes, the system 200 may determine which index to access based on rules defined by a database architect or system administrator. For instance, a database architect or system administrator may set rules defining data of interest to an organization and the system 200 may access the index corresponding to the data of interest.

In some implementations, the system 200 may determine which index to access dynamically based on the user or device requesting access. In these implementations, the system 200 may determine access level credentials of the user or device that submitted the query and determine which index to access based on the determined credentials. In addition, the rules may define that different users or different types of users receive different indexes. For instance, the system 200 may access an index for financial data of an organization when the user accessing the data is a financial analyst, but may access

25

an index for personnel data of the organization when the user accessing the data is a human resources manager.

In some examples, the system **200** may generate the index based on receipt of the query. In these example, the data of interest may be compressed and organized prior to receiving the query and the system **200** may generate an index that is descriptive of the compressed and organized data based on the received query. In this regard, the system **200** may identify the compressed and organized data (e.g., attribute classes) of interest and dynamically generate an index for only the data (e.g., attribute classes) of interest.

The system **200** identifies relative positions in the accessed index for one or more attribute classes included in the received query (**1030**). The system **200** may analyze the accessed index to identify relationships between elements in the index. The index may define parent-child relationships between different elements or portions of the index. The elements or portions of the index may be nodes in a graph structure or blocks in an element index. The index may store data defining relationships between the elements in the index and the data may be used to identify parent-child relationships between elements in the index. Throughout the disclosure, the terms higher and lower are used to refer to relative positions of elements within the index, with a higher level attribute class having a parent relationship to a lower level attribute class and the lower level attribute class having a child relationship to the higher level attribute class. The system **200** may analyze the index, identify the parent-child relationships between the attribute classes in the query, and determine relative positions of the attribute classes based on the identified parent-child relationships.

To illustrate, suppose the system **200** accessed the index shown in FIGS. 7 and 8. The index shown in FIGS. 7 and 8 includes the attribute classes “year,” “month,” “state,” and “city,” with the “year” attribute class being an immediate parent of the “month” attribute class, the “month” attribute class being an immediate parent of the “state” attribute class, and the “state” attribute class being an immediate parent of the “city” attribute class. Accordingly, in this example, the system **200** determines that the “year” attribute class has a higher relative position than each of the “month” attribute class, the “state” attribute class, and the “city” attribute class. The “month” attribute class has a lower relative position than the “year” attribute class, but a higher relative position than each the “state” attribute class and the “city” attribute class. The “state” attribute class has a lower relative position than each of the “year” attribute class and the “month” attribute class, but a higher relative position than the “city” attribute class. The “city” attribute class has a lower relative position than each of the “year” attribute class, the “month” attribute class, and the “state” attribute class. Based on the query shown in the input interface **1010** and using the index shown in FIGS. 7 and 8, the system **200** determines that the “year” attribute class has a higher relative position (e.g., the highest relative position of relevant attribute classes) than the state attribute class (e.g., the lowest relative position of relevant attribute classes).

The system **200** processes the accessed index based on the identified relative positions (**1040**). For instance, the system **200** may analyze the index in light of the relative positions of the attribute classes associated with the received query. The system **200** may process the index starting with the portion of the index corresponding to the attribute class with the highest relative position in the index and continue processing the index in descending order to the attribute class with the lowest relative position in the index. This type of processing is described in more detail below with respect to FIGS. 11-15.

26

In addition, the system **200** may process the index by identifying the attribute class with the lowest relative position and using the addressing information associated with the portions of the index having the lowest relative position to group data records. In this regard, the system **200** may start analyzing the attribute class with the lowest relative position and identify values in the attribute class with the lowest relative position that have unique values and combinations of parent values. The system **200** further may use the identified relative positions in identifying redundancies within the index as described in more detail below with respect to FIG. 16.

The system **200** groups, using the processed index, data records with common values for the one or more attribute classes included in the received query (**1050**). For example, the system **200** may compare values for the relevant attribute classes in the index and identify all unique combinations of values for the relevant attribute classes based on the comparison. In this example, the system **200** may identify groups of data records by identifying the portions of the index having each of the unique combinations of attribute values and accessing address or storage location information for the identified portions. The accessed address or storage location information corresponds to the data records included in the group and may be used to access the data records in the group. Techniques for grouping data records using the processed index are described in more detail below with respect to FIGS. 11-18.

As one example, suppose a user enters the query shown in the input interface **1010** and the system **200** uses the index shown in FIGS. 7 and 8. In this example, the system **200** analyzes the index shown in FIGS. 7 and 8 and determines that five unique combinations of year and state values exist: “06-VA,” “06-NY,” “07-NY,” “07-MD,” and “07-VA.” The system **200** then groups the portions of the index together that correspond to each of the identified combinations. In this regard, the system **200** may process the index **800** and determine that the “06-VA” combination corresponds to rows 0 and 2-3 in the database table, the “06-NY” combination corresponds to row 1 in the database table, the “07-NY” combination corresponds to row 4 in the database table, the “07-MD” combination corresponds to row 5 in the database table, and the “07-VA” combination corresponds to row 6 in the database table.

For each group of data records, the system **200** accesses metrics associated with the group of data records and performs computations on the accessed metrics (**1060**). The system **200** may automatically, without user intervention, generate database code for accessing data corresponding to the identified groups of data records and access the metrics using the generated database code. The system **200** may generate the database code using the address or storage location information included in the index and associated with each identified group. For instance, the system **200** may identify addresses of data records mapped to by the processed index and access, from electronic data storage, metrics corresponding to the identified addresses of data records mapped to by the processed index. The system **200** performs a computation for each group of the accessed metrics to use in the report.

In the example discussed above in which a user enters the query shown in the input interface **1010** and the system **200** uses the index shown in FIGS. 7 and 8, the system **200** uses the identified database row information to access the relevant information for each unique combination of values identified for the year and state attribute classes. For the “06-VA” combination, the system **200** accesses the sales metric stored in each of the rows 0 and 2-3 in the database table (e.g., values 0.8, 1.7, and 1.4 in the example data shown in FIG. 6) and

computes an average value for the accessed sales metrics (e.g., 1.3 in the example data shown in FIG. 6). For the “06-NY” combination, the system 200 accesses the sales metric stored in row 1 in the database table (e.g., value 2.0 in the example data shown in FIG. 6) and, because it is the only metric for the combination, uses that value as the average sales metric. For the “07-NY” combination, the system 200 accesses the sales metric stored in row 4 in the database table (e.g., value 2.2 in the example data shown in FIG. 6) and, because it is the only metric for the combination, uses that value as the average sales metric. For the “07-MD” combination, the system 200 accesses the sales metric stored in row 5 in the database table (e.g., value 1.6 in the example data shown in FIG. 6) and, because it is the only metric for the combination, uses that value as the average sales metric. For the “07-VA” combination, the system 200 accesses the sales metric stored in row 6 in the database table (e.g., value 1.2 in the example data shown in FIG. 6) and, because it is the only metric for the combination, uses that value as the average sales metric.

The system 200 generates a report 1080 based on the computations performed on the accessed metrics (1070). For instance, the system 200 may generate a report that includes the result of the computation for each group of the accessed metrics along with values for the multiple attribute classes identified as being related to the input query that correspond to the group of the accessed metrics. The report may include information identifying the type of report generated, computations made on the accessed metrics, and identifiers associated with the computed metrics that identify to which group of attribute classes the computed metric belongs. The report may be stored in electronic storage, displayed on any type of display device, and/or communicated over a network to another device.

As shown, the report 1080 corresponds to the example discussed above in which a user enters the query shown in the input interface 1010 and the system 200 uses the index shown in FIGS. 7 and 8. The report 1080 includes graphical identifiers that indicate each of the five unique combinations of values for the relevant attribute classes and the average sale value metric computed for each of the five unique combinations of values. Other types of reports may be generated and the reports may include charts, graphs, tables, and any other type of graphical format for displaying a representation of data. In addition, the system 200 may generate reports that include multiple computed metrics.

FIG. 11 illustrates a process 1100 for grouping data records. The operations of the process 1100 are described generally as being performed by the system 200. The operations of the process 1100 may be performed exclusively by the data processing system 205, may be performed exclusively by the database system 280, or may be performed by a combination of the data processing system 205 and the database system 280. In some implementations, operations of the process 1100 may be performed by one or more processors included in one or more electronic devices.

The system 200 removes portions of an index that correspond to attribute classes that are not associated with a received query (1110). For instance, the system 200 identifies attribute classes included in a received query and updates an accessed index by removing portions or elements (e.g., nodes or blocks) of the index that correspond to an attribute class other than the identified attribute classes. Based on the removal, relationships between the remaining portions or elements in the index may be incomplete. As such, the system 200 may update the relationships defined by the index to

reflect the parent-child relationships and relative positions of the remaining portions or elements in the index.

FIG. 12 illustrates an example of a graph data structure 1200 in which a single attribute class has been processed in accordance with the process 1100. The graph data structure 1200 is based on processing of the index shown in FIGS. 7 and 8. In this example, the system 200 has received a query listing the attribute classes “year,” “state,” and “city.” Because the query does not include the attribute class “month,” the system 200 removes portions or elements in the index that correspond to the attribute class “month” and updates the relationships defined in the index based on the removal. As reflected in a comparison of the graph data structure 1200 to the graph data structure 700 shown in FIG. 7, the system 200 has removed nodes 720, 730, and 750 from the graph and has updated the connections within the graph to show a direct parent-child relationship between the nodes associated with the year attribute class and the nodes associated with the state attribute class.

Referring again to FIG. 11, the system 200 identifies relative positions in the index for the attribute classes associated with the received query (1120). For example, the system 200 may analyze the index after removal of the portions of the index that correspond to attribute classes that are not associated with a received query and identify parent-child relationships between the attribute classes remaining in the index. In the example shown in FIG. 12, the system 200 determines that the year attribute class has the highest relative position within the index, the state attribute class has the next highest relative position in the index, and the city attribute class has the lowest relative position in the index.

The system 200 selects the attribute class associated with the received query that has the highest relative position in the index (1130). For instance, the system 200 may compare the identified relative positions and select the attribute class having the highest relative position for the attribute classes remaining in the index. In the example shown in FIG. 12, the system 200 determines that the year attribute class has the highest relative position within the index.

The system 200 identifies a first distinct value within the selected attribute class (1140). For example, the system 200 accesses the first portion or element (e.g., node or block) of the index within the selected attribute class and retrieves the data value associated with the first portion or element of the index. Because the data value is associated with the first portion or element of the index, it is the first distinct value within the selected attribute class. In the example shown in FIG. 12, the system 200 identifies the first distinct value within the selected attribute class as being the value “06.” The system 200 may store the first distinct value for comparison against other values within the selected attribute class.

The system 200 selects a distinct identifier to assign to portions of the index within the selected attribute class with the identified distinct value (1150). The distinct identifier may be any type of identifier used to denote and identify a portion or element of an index uniquely. For instance, the distinct identifier may be an alphanumeric character that labels the portion or element of the index. The system 200 may use a running count of numeric values as the distinct identifiers, with the running count being incremented each time a new distinct value is identified and processed. In the example shown in FIG. 12, the system 200 selects the identifier “0” as the first distinct identifier to use in labeling nodes associated with the value “06” in the year attribute class in the graph data structure 1200.

The system 200 assigns the selected identifier to all portions of the index within the selected attribute class with the

29

identified distinct value (1160). For instance, the system 200 may compare the values of each of the nodes within the selected attribute class to the first distinct value and store data that labels each matching node with the distinct identifier. In the example shown in FIG. 12, because the value "06" is found in only the node 1210, the system 200 labels the node 1210 using the distinct identifier "0."

The system 200 determines whether more distinct values within the selected attribute class exist (1170). The system 200 may determine whether all of the nodes within the year attribute class have already been processed by tracking processing of the index and recognizing when the last portion or element of the index has been processed. The system 200 also may determine whether any portions or elements have not been assigned an identifier and, thereby, determine that more distinct values exist when portions or elements within the selected attribute class have not been assigned an identifier and determine that more distinct values do not exist when all portions or elements within the selected attribute class have been assigned an identifier. In the example shown in FIG. 12, because the value "07" exists and the node 1250 has not been assigned an identifier at this point, the system 200 determines that more distinct data values exist.

When the system 200 determines that more distinct values exist, the system 200 identifies a next distinct value within the selected attribute class (1175). For example, the system 200 accesses the next (e.g., second) portion or element (e.g., node or block) of the index within the selected attribute class and retrieves the data value associated with the next portion or element of the index. The system 200 may analyze the portions or elements in the index in order to identify the next distinct value within the selected attribute class. The system 200 may analyze the portions or elements in the index that have the same relative position in the index as the selected attribute class. In the example shown in FIG. 12, the system 200 identifies the next distinct value within the selected attribute class as being the value "07." The system 200 may store the next distinct value for comparison against other values within the selected attribute class.

After identifying the next distinct value within the selected attribute class, the system 200 repeats the operations described above with respect to reference numerals 1150, 1160, and 1170. In the example shown in FIG. 12, for the next distinct value "07," the system 200 selects the distinct identifier "1" and assigns the distinct identifier to node 1250, which is associated with the value "07." At this point, the system 200 determines that more distinct data values within the selected attribute class do not exist.

When the system 200 determines that more distinct values do not exist, the system 200 determines whether more attribute classes exist (1180). The system 200 may determine whether all of the attribute classes associated with the received query have been processed.

When the system 200 determines that more attribute classes exist, the system 200 processes the additional attribute classes (1185). Processing the additional attribute classes is described in more detail below with respect to FIG. 13.

When the system 200 determines that more attribute classes do not exist, the system 200 groups data records referred to by portions of the index within the selected attribute class and assigned the same identifier (1190). For example, the system 200 may compare the identifiers assigned to the portions of the index within the selected attribute class with one another and identify portions of the index that have the same identifier. In the example shown in FIG. 12, if additional attribute classes did not exist after processing the year attribute class (e.g., the only attribute

30

class associated with the received query was the year attribute class), the system 200 would identify two groups of data records. Specifically, a first group of data records would be associated with the node 1210 and a second group of data records would be associated with the node 1250. To process the first group of data records associated with the node 1210, the system 200 may access storage location mapping information associated with the node 1210 in the index and access the first group of data records using the storage location mapping information. To process the second group of data records associated with the node 1250, the system 200 may access storage location mapping information associated with the node 1250 in the index and access the second group of data records using the storage location mapping information.

FIG. 13 illustrates a process 1300 for grouping data records. The operations of the process 1300 are described generally as being performed by the system 200. The operations of the process 1300 may be performed exclusively by the data processing system 205, may be performed exclusively by the database system 280, or may be performed by a combination of the data processing system 205 and the database system 280. In some implementations, operations of the process 1300 may be performed by one or more processors included in one or more electronic devices.

The system 200 selects the attribute class associated with the received query that has the next highest relative position in the index (1310). For instance, the system 200 may compare the identified relative positions of the relevant attribute classes within the index and select the attribute class having the next highest relative position of the attribute classes remaining in the index. When the system 200 has only processed the attribute class with the highest relative position, the system 200 selects the attribute class with the second highest relative position as the immediate child of the attribute class with the highest relative position.

FIG. 14 illustrates an example of a graph data structure 1400 in which two attribute classes have been processed in accordance with the process 1100 and the process 1300. The graph data structure 1400 corresponds to the graph data structure 1200, except that it is used to show additional processing of the attribute class having the second highest relative position in the index. In the example shown in FIG. 14, the system 200 determines that the state attribute class has the next highest relative position within the index at this point in the analysis.

The system 200 identifies a first distinct pair of distinct value within the selected attribute class and distinct identifier assigned to immediate parent in the index (1320). For example, the system 200 accesses the first portion or element (e.g., node or block) of the index within the selected attribute class and retrieves the data value associated with the first portion or element of the index within the selected attribute class. Because the data value is associated with the first portion or element of the index within the selected attribute class, it is the first distinct value within the selected attribute class. In addition, the system 200 identifies the immediate parent in the index for the first portion or element of the index within the selected attribute class and accesses the identifier assigned to the immediate parent in prior processing. The first distinct value within the selected attribute class and the identifier of its immediate parent form the first distinct pair. In the example shown in FIG. 14, the system 200 identifies the first distinct value within the selected attribute class as being the value "VA" associated with the node 1420 and the identifier for its immediate parent as being the identifier "0" assigned to the

31

node **1210**. The system **200** may store the first distinct pair for comparison against other pairs within the selected attribute class.

The system **200** selects a distinct identifier to assign to portions of the index within the selected attribute class with the identified distinct pair (**1330**). The distinct identifier may be any type of identifier used to denote and identify a portion or element of an index uniquely. For instance, the distinct identifier may be an alphanumeric character that labels the portion or element of the index. The system **200** may use a running count of numeric values as the distinct identifiers, with the running count being incremented each time a new distinct value is identified and processed. In some implementations, the distinct identifier may be distinct within the selected attribute class, but not necessarily distinct with respect to other attribute classes in the index. In the example shown in FIG. **14**, the system **200** selects the identifier "0" as the first distinct identifier to use in labeling nodes associated with the first distinct pair (e.g., "VA-0") in the graph data structure **1400**.

The system **200** assigns the selected identifier to all portions of the index within the selected attribute class with the identified distinct pair (**1340**). For instance, the system **200** may compare the values of each of the nodes within the selected attribute class to the first distinct value and, for matching values, compare the identifier in the distinct pair to the identifier assigned to the immediate parent of the node with the matching value. When both the value of the node within the selected attribute class and the identifier assigned to its immediate parent match the identified distinct pair, the system **200** may store data that labels the corresponding node with the distinct identifier.

In the example shown in FIG. **14**, the system **200** labels the node **1420** using the distinct identifier "0" and determines that nodes **1440** and **1480** have a value (e.g., "VA") matching the value in the identified distinct pair (e.g., "VA-0") while nodes **1430**, **1460**, and **1470** have values that do not match. Accordingly, the system **200** accesses the identifier assigned to the immediate parent node for the nodes **1440** and **1480** and compares the assigned identifier to the identifier included in the distinct pair (e.g., "0"). For the node **1440**, the system **200** accesses the identifier assigned to node **1210** (e.g., "0"), determines that node **1440** matches the identified distinct pair, and, therefore, labels the node **1440** using the distinct identifier "0." For the node **1480**, the system **200** accesses the identifier assigned to node **1250** (e.g., "1"), determines that node **1480** does not match the identified distinct pair, and, therefore, does not label the node **1480** using the distinct identifier "0," despite the node **1480** having a value that matches the value in the identified distinct pair.

The system **200** determines whether more distinct pairs exist (**1350**). The system **200** may determine whether all of the nodes within the selected attribute class have already been processed by tracking processing of the index and recognizing when the last portion or element of the index has been processed. The system **200** also may determine whether any portions or elements within the selected attribute class have not been assigned an identifier and, thereby, determine that more distinct pairs exist when portions or elements within the selected attribute class have not been assigned an identifier and determine that more distinct pairs do not exist when all portions or elements within the selected attribute class have been assigned an identifier. In the example shown in FIG. **14**, several nodes have not been assigned an identifier at this point and, therefore, the system **200** determines that more distinct pairs exist.

32

When the system **200** determines that more distinct pairs exist, the system **200** identifies a next distinct pair within the selected attribute class (**1360**). For example, the system **200** identifies the next portion or element (e.g., node or block) of the index within the selected attribute class that has not been assigned an identifier. The system **200** then retrieves the data value associated with the identified next portion or element of the index and the identifier assigned to its immediate parent. In the example shown in FIG. **14**, the system **200** identifies the next distinct pair within the selected attribute class as being the value "NY" and the parent identifier "0." The system **200** may store the next distinct pair for comparison against other values within the selected attribute class.

After identifying the next distinct pair within the selected attribute class, the system **200** repeats the operations described above with respect to reference numerals **1340** and **1350**. In the example shown in FIG. **14**, for the next distinct pair "NY-0," the system **200** selects the distinct identifier "1" and assigns the distinct identifier to node **1430**, which is the only node in the selected attribute class associated with the identified distinct pair of "NY-0." At this point, the system **200** determines that more distinct pairs within the selected attribute class exist and identifies the next distinct pair within the selected attribute class as being the value "NY" and the parent identifier "1." For the next distinct pair "NY-1," the system **200** selects the distinct identifier "2" and assigns the distinct identifier to node **1460**, which is the only node in the selected attribute class associated with the identified distinct pair of "NY-1." At this point, the system **200** determines that more distinct pairs within the selected attribute class still exist and identifies the next distinct pair within the selected attribute class as being the value "MD" and the parent identifier "1." For the next distinct pair "MD-1," the system **200** selects the distinct identifier "3" and assigns the distinct identifier to node **1470**, which is the only node in the selected attribute class associated with the identified distinct pair of "MD-1." At this point, the system **200** determines that more distinct pairs within the selected attribute class still exist and identifies the next distinct pair within the selected attribute class as being the value "VA" and the parent identifier "1." For the next distinct pair "VA-1," the system **200** selects the distinct identifier "4" and assigns the distinct identifier to node **1480**, which is the only node in the selected attribute class associated with the identified distinct pair of "VA-1." At this point, the system **200** determines that more distinct pairs within the selected attribute class do not exist because all nodes within the selected attribute class have been assigned an identifier.

When the system **200** determines that more distinct values do not exist, the system **200** determines whether more attribute classes exist (**1370**). The system **200** may determine whether all of the attribute classes associated with the received query have been processed.

When the system **200** determines that more attribute classes exist, the system **200** continues processing and repeats the operations described above with respect to reference numerals **1310**, **1320**, **1330**, **1340**, **1350**, **1360**, and **1370** until all attribute classes have been processed and all nodes within the index have been assigned an identifier.

FIG. **15** illustrates an example of a graph data structure **1500** in which three attribute classes have been processed in accordance with the process **1100** and the process **1300**. The graph data structure **1500** corresponds to the graph data structure **1200** and the graph data structure **1400**, except that it is used to show additional processing of the attribute class having the third highest relative position in the index. In the example shown in FIG. **15**, the system **200** determines that the

city attribute class has the next highest relative position within the index at this point in the analysis.

After determining that the city attribute class has the next highest relative position within the index, the system 200 identifies the first distinct pair of distinct value within the selected attribute class and distinct identifier assigned to immediate parent in the index as the value "McLean" associated with the node 1522 and the identifier "0." The system 200 selects the identifier "0" as the first distinct identifier to use in labeling nodes in the selected attribute class associated with the first distinct pair (e.g., "McLean-0") in the graph data structure 1500. Then, the system 200 labels the node 1522 using the distinct identifier "0" and determines that node 1542 has a value (e.g., "McLean") matching the value in the identified distinct pair (e.g., "McLean-0") while nodes 1532, 1544, 1562, 1572, and 1582 have values that do not match. Accordingly, the system 200 accesses the identifier assigned to the immediate parent node for the node 1542 and compares the assigned identifier to the identifier included in the distinct pair (e.g., "0"). For the node 1542, the system 200 accesses the identifier assigned to node 1440 (e.g., "0"), determines that node 1542 matches the identified distinct pair, and, therefore, labels the node 1542 using the distinct identifier "0."

The system 200 determines that more distinct pairs within the selected attribute class exist and identifies the next distinct pair within the selected attribute class as being the value "New York" and the parent identifier "1." For the next distinct pair "New York-1," the system 200 selects the distinct identifier "1" and assigns the distinct identifier to node 1532, which is the only node in the selected attribute class associated with the identified distinct pair of "New York-1." At this point, the system 200 determines that more distinct pairs within the selected attribute class exist and identifies the next distinct pair within the selected attribute class as being the value "Vienna" and the parent identifier "0." For the next distinct pair "Vienna-0," the system 200 selects the distinct identifier "2" and assigns the distinct identifier to node 1544, which is the only node in the selected attribute class associated with the identified distinct pair of "Vienna-0." At this point, the system 200 determines that more distinct pairs within the selected attribute class still exist and identifies the next distinct pair within the selected attribute class as being the value "New York" and the parent identifier "2." For the next distinct pair "New York-2," the system 200 selects the distinct identifier "3" and assigns the distinct identifier to node 1562, which is the only node in the selected attribute class associated with the identified distinct pair of "New York-2." At this point, the system 200 determines that more distinct pairs within the selected attribute class still exist and identifies the next distinct pair within the selected attribute class as being the value "Baltimore" and the parent identifier "3." For the next distinct pair "Baltimore-3," the system 200 selects the distinct identifier "4" and assigns the distinct identifier to node 1572, which is the only node in the selected attribute class associated with the identified distinct pair of "Baltimore-3." At this point, the system 200 determines that more distinct pairs within the selected attribute class still exist and identifies the next distinct pair within the selected attribute class as being the value "Vienna" and the parent identifier "4." For the next distinct pair "Vienna-4," the system 200 selects the distinct identifier "5" and assigns the distinct identifier to node 1582, which is the only node in the selected attribute class associated with the identified distinct pair of "Vienna-4." At this point, the system 200 determines that more distinct pairs within the selected attribute class do not exist because all nodes within the selected attribute class have been assigned an identifier. The system 200 also deter-

mines that more attribute classes do not exist because all of the attribute classes have been processed and all nodes have been assigned an identifier.

When the system 200 determines that more attribute classes do not exist, the system 200 groups data records referred to by portions of the index within the attribute class with the lowest relative position and assigned the same identifier (1380). For example, the system 200 may compare the identifiers assigned to the portions of the index within the attribute class having the lowest relative position within the index with one another and identify portions of the index that have the same identifier. The system 200 groups the portions of the index within the attribute class having the lowest relative position within the index and having the same identifier together and accesses storage location information for data records associated with the groups.

In the example shown in FIG. 15, the system 200 identifies the city attribute class (e.g., nodes 1522, 1532, 1542, 1544, 1562, 1572, and 1582) as the attribute class having the lowest relative position within the index. Accordingly, the system 200 analyzes the identifiers assigned to the nodes in the city attribute class (e.g., nodes 1522, 1532, 1542, 1544, 1562, 1572, and 1582) to identify groups of nodes having the same identifier. In this example, the nodes 1522 and 1542 are the only nodes that share the same identifier and are grouped together. The other nodes 1532, 1544, 1562, 1572, and 1582 are each identified as a group including a single node.

Based on the groupings, the system 200 accesses storage location mapping information associated with each group of nodes. The storage location mapping information associated with the graph data structure 1500 may be the element index 800 shown in FIG. 8. In this case, the system 200 uses the element index 200 to identify rows in the database table that correspond to the groups of nodes identified in the graph data structure 1500. Because the groups of nodes correspond to the city attribute class, the system 200 uses the table 880 to map the groups of nodes to rows in the database table. For the group including nodes 1522 and 1542, the system 200 determines that these nodes corresponds to blocks 0 and 2 in the city attribute class, respectively. As such, the system 200 identifies entries in the table 880 that correspond to blocks 0 and 2 and retrieves the row information stored in the identified entries. In this example, the system 200 determines the row 0 and row 2 in the database table correspond to blocks 0 and 2 and groups those rows together as data records sharing a unique combination of values for the attribute classes included in the received query. The other groups of nodes are processed in a similar manner and the system 200 is able to determine groups of one or more rows in the database table for each unique combination of values for the attribute classes included in the received query.

FIG. 16 illustrates another process 1600 for grouping data records. The operations of the process 1600 are described generally as being performed by the system 200. The operations of the process 1600 may be performed exclusively by the data processing system 205, may be performed exclusively by the database system 280, or may be performed by a combination of the data processing system 205 and the database system 280. In some implementations, operations of the process 1600 may be performed by one or more processors included in one or more electronic devices.

The system 200 modifies an index to remove portions of the index that correspond to attribute classes that are not associated with the received query (1610). The system 200 may modify an index to remove portions of the index that correspond to attribute classes that are not associated with the

received query using techniques similar to those discussed above with respect to reference numeral **1110**.

The system **200** analyzes the modified index (**1620**) and identifies redundancies present in the modified index (**1630**). For instance, the system **200** may compare values of portions of the index within the same attribute class and determine whether any portions include the same value and the same parent in the index. These types of redundancies may occur because of the removal of portions of the index. Specifically, two portions of an index may share a common value, a common grandparent in the index, but a different parent in the index. These two portions are not redundant because of the difference in parents. However, when the parent attribute class is removed, the two portions become redundant because they share the same value and the same grandparent.

The system **200** compresses the identified redundancies in the modified index (**1640**). For example, the system **200** may modify the index to eliminate redundant portions or elements (e.g., redundant blocks or nodes). In this example, the system **200** may collapse two portions or elements in the index identified as being redundant into a single portion or element in the index.

FIG. **17** illustrates an example of a graph data structure **1700** in which an index has been processed in accordance with the process **1600**. The graph data structure includes nodes **1710** and **1740**, which correspond to the year attribute class, nodes **1720**, **1730**, **1750**, **1760**, and **1770**, which correspond to the state attribute class, and nodes **1722**, **1732**, **1734**, **1752**, **1762**, and **1772**, which correspond to the city attribute class.

The graph data structure **1700** is based on processing of the index shown in FIGS. **7** and **8**. In this example, the system **200** has received a query listing the attribute classes “year,” “state,” and “city.” Because the query does not include the attribute class “month,” the system **200** removes portions or elements in the index that correspond to the attribute class “month” and updates the relationships defined in the index based on the removal. As reflected in a comparison of the graph data structure **1700** to the graph data structure **700** shown in FIG. **7**, the system **200** has removed nodes **720**, **730**, and **750** from the graph and has updated the connections within the graph to show a direct parent-child relationship between the nodes associated with the year attribute class and the nodes associated with the state attribute class.

In addition, the system **200** has identified and compressed redundancies resulting from the removal of the month attribute class. The graph data structure **1500** illustrates an example in which redundancies resulting from the removal of the month attribute class have not been compressed. As shown in FIG. **15**, the nodes **1420** and **1440** have become redundant based on removal of the month attribute class. In preparing the graph data structure **1700**, the system **200** identifies this redundancy and collapses the nodes **1420** and **1440** into a single node **1730**. Also, as shown in FIG. **15**, the nodes **1522** and **1542** have become redundant based on removal of the month attribute class. In preparing the graph data structure **1700**, the system **200** identifies this redundancy and collapses the nodes **1522** and **1542** into a single node **1732**.

Referring again to FIG. **16**, the system **200** updates index mapping definitions based on compression of the identified redundancies in the modified index (**1650**). Because each portion of the index that is compressed is associated with its own set of storage location mapping definitions, the system **200** may need to associate the mapping definitions for each compressed portion with the portion of the index that remains after compression. For instance, when two redundant portions of an index are compressed into a single portion, the

system **200** may identify the mapping definitions associated with the first redundant portion, identify the mapping definitions associated with the second redundant portion, and associate the single compressed portion of the index with both the identified mapping definitions associated with the first redundant portion and the identified mapping definitions associated with the second redundant portion.

FIG. **18** illustrates an example of an updated element index in which the element index has been processed in accordance with the process **1600**. The element index includes a year column **1810**, a state column **1820**, and a city column **1830**. For each attribute class, the element index includes one entry for each distinct data value within the corresponding attribute class. For each distinct data value, the element index **800** includes the distinct data value and block identification data that identifies which one or more blocks within the attribute class have the corresponding data value.

The element index also is associated with tables that correspond to attribute classes included in the accessed data. The tables store a block identifier for each block included in an attribute class and store location information for data records that are included in the corresponding block. As shown, the table **1840** corresponds the year attribute class. The table **1840** includes an entry **1841** corresponding to the first block (e.g., Block **0**) in the year attribute class and entry **1842** corresponding to the second block (e.g., Block **1**) in the year attribute class. The entries **1841-1842** define which rows in a database table correspond to the block identified in the corresponding entry and the table **1840** includes data associating each row in the database with a block.

The table **1850** corresponds the state attribute class. The table **1850** includes an entry **1851** corresponding to the first block (e.g., Block **0**) in the state attribute class, an entry **1852** corresponding to the second block (e.g., Block **1**) in the state attribute class, an entry **1853** corresponding to the third block (e.g., Block **2**) in the state attribute class, an entry **1854** corresponding to the fourth block (e.g., Block **3**) in the state attribute class, and an entry **1855** corresponding to the fifth block (e.g., Block **4**) in the state attribute class. The entries **1851-1855** define which rows in a database table correspond to the block identified in the corresponding entry and the table **1850** includes data associating each row in the database with a block.

The table **1860** corresponds the city attribute class. The table **1860** includes an entry **1861** corresponding to the first block (e.g., Block **0**) in the city attribute class, an entry **1862** corresponding to the second block (e.g., Block **1**) in the city attribute class, an entry **1863** corresponding to the third block (e.g., Block **2**) in the city attribute class, an entry **1864** corresponding to the fourth block (e.g., Block **3**) in the city attribute class, an entry **1865** corresponding to the fifth block (e.g., Block **4**) in the city attribute class, and an entry **1866** corresponding to the sixth block (e.g., Block **5**) in the city attribute class. The entries **1861-1866** define which rows in a database table correspond to the block identified in the corresponding entry and the table **1860** includes data associating each row in the database with a block.

The element index shown in FIG. **18** is based on processing of the index shown in FIGS. **7** and **8** and corresponds to the graph data structure **1700** shown in FIG. **17**. In this example, the system **200** has received a query listing the attribute classes “year,” “state,” and “city.” Because the query does not include the attribute class “month,” the system **200** removes portions or elements in the index that correspond to the attribute class “month.” As reflected in a comparison with the element index **800** shown in FIG. **8**, the system **200** has



37

removed the month column **820** and the month table **860** because the month attribute class was not included in the received query.

In addition, the system **200** has updated the mapping definitions in the element index based on compression of redundant portions of the index. For example, the system **200** identified the blocks **0** and **2** in the state attribute class shown in FIG. **8** as being redundant and collapsed them into a single block. Specifically, the system **200** removed the block **0** and collapsed it into block **2**. Based on this compression operation, the system **200** updated the element index by first removing block **0** in the state attribute class and renumbering the remaining blocks (e.g., changed block **1** to block **0**, changed block **2** to block **1**, etc.). After renumbering the blocks, the system **200** updated the mapping definitions to reflect the compression of block **0** into block **2** (now renumbered as block **1**) by associating the mapping definitions for the two redundant blocks with the single compressed block. As shown, the entry **1852** associates block **1** (i.e., previously block **2** prior to renumbering) with both the mapping definitions from block **0** and block **2** (e.g., row **0** and rows **2** to **3**) of the state attribute class shown in FIG. **8**. Accordingly, the mapping definition of the element index has been preserved after the compression.

Similarly, the system **200** identified the blocks **0** and **2** in the city attribute class shown in FIG. **8** as being redundant and collapsed them into a single block. Specifically, the system **200** removed the block **0** and collapsed it into block **2**. Based on this compression operation, the system **200** updated the element index by first removing block **0** in the city attribute class and renumbering the remaining blocks (e.g., changed block **1** to block **0**, changed block **2** to block **1**, etc.). After renumbering the blocks, the system **200** updated the mapping definitions to reflect the compression of block **0** into block **2** (now renumbered as block **1**) by associating the mapping definitions for the two redundant blocks with the single compressed block. As shown, the entry **1862** associates block **1** (i.e., previously block **2** prior to renumbering) with both the mapping definitions from block **0** and block **2** (e.g., row **0** and row **2**) of the city attribute class shown in FIG. **8**. Accordingly, the mapping definition of the element index has been preserved after the compression.

Referring again to FIG. **16**, the system **200** groups data records based on the updated index (**1660**). For instance, the system **200** identifies portions of the index associated with the attribute class having the lowest relative position in the updated index. For each portion, the system **200** identifies storage location mapping definitions and uses those storage location mapping definitions as a group of data records. Because the index has been updated to compress redundancies and update storage location mapping definitions, each portion or element in the updated index that is associated with the attribute class having the lowest relative position in the index is associated with a group of data records that has a unique combination of values for the attribute classes included in the received query. As such, after updating the index, the system **200** may only need to identify data records mapped to by a portion of the index associated with the attribute class having the lowest relative position in the index to identify a group data records having a unique combination of values for the attribute classes included in the received query.

FIG. **19** is a schematic diagram of a generic computer system **1900**. The system **1900** can be used for the operations described in association with any of the computer-implemented methods described previously. The system **1900** includes a processor **1910**, a memory **1920**, a storage device

38

**1930**, and an input/output device **1940**. Each of the components **1910**, **1920**, **1930**, and **1940** are interconnected using a system bus **1950**. The processor **1910** is capable of processing instructions for execution within the system **1900**. In one implementation, the processor **1910** is a single-threaded processor. In another implementation, the processor **1910** is a multi-threaded processor. The processor **1910** is capable of processing instructions stored in the memory **1920** or on the storage device **1930** to display graphical information for a user interface on the input/output device **1940**.

The memory **1920** stores information within the system **1900**. In some implementations, the memory **1920** is a computer-readable storage medium. In another implementation, the memory **1920** is a volatile memory unit. In yet another implementation, the memory **1920** is a non-volatile memory unit.

The storage device **1930** is capable of providing mass storage for the system **1900**. In some implementation, the storage device **1930** is a computer-readable storage medium. In various different implementations, the storage device **1930** may be a floppy disk device, a hard disk device, an optical disk device, or a tape device.

The input/output device **1940** provides input/output operations for the system **1900**. In one implementation, the input/output device **1940** includes a keyboard and/or pointing device. In another implementation, the input/output device **1940** includes a display unit for displaying graphical user interfaces.

The features described can be implemented in digital electronic circuitry, or in computer hardware, firmware, software, or in combinations of them. An apparatus can be implemented in a computer program product tangibly embodied in a machine-readable storage device, for execution by a programmable processor; and method steps can be performed by a programmable processor executing a program of instructions to perform functions of the described implementations by operating on input data and generating output. The described features can be implemented in one or more computer programs that are executable on a programmable system including at least one programmable processor coupled to receive data and instructions from, and to transmit data and instructions to, a data storage system, at least one input device, and at least one output device. A computer program is a set of instructions that can be used, directly or indirectly, in a computer to perform a certain activity or bring about a certain result. A computer program can be written in any form of programming language, including compiled or interpreted languages, and it can be deployed in any form, including as a stand-alone program or as a module, component, subroutine, or other unit suitable for use in a computing environment.

Suitable processors for the execution of a program of instructions include, by way of example, both general and special purpose microprocessors, and a sole processor or one of multiple processors of any kind of computer. Generally, a processor will receive instructions and data from a read-only memory or a random access memory or both. The elements of a computer are a processor for executing instructions and one or more memories for storing instructions and data. Generally, a computer will also include, or be operatively coupled to communicate with, one or more mass storage devices for storing data files; such devices include magnetic disks, such as internal hard disks and removable disks; magneto-optical disks; and optical disks. Storage devices suitable for tangibly embodying computer program instructions and data include all forms of non-volatile memory, including by way of example semiconductor memory devices, such as EPROM, EEPROM, and flash memory devices; magnetic disks such as



39

internal hard disks and removable disks; magneto-optical disks; and CD-ROM and DVD-ROM disks. The processor and the memory can be supplemented by, or incorporated in, ASICs (application-specific integrated circuits).

To provide for interaction with a user, the features can be implemented on a computer having a display device such as a CRT (cathode ray tube) or LCD (liquid crystal display) monitor for displaying information to the user and a keyboard and a pointing device such as a mouse or a trackball by which the user can provide input to the computer.

The features can be implemented in a computer system that includes a back-end component, such as a data server, or that includes a middleware component, such as an application server or an Internet server, or that includes a front-end component, such as a client computer having a graphical user interface or an Internet browser, or any combination of them. The components of the system can be connected by any form or medium of digital data communication such as a communication network. Examples of communication networks include, e.g., a LAN, a WAN, and the computers and networks forming the Internet.

The computer system can include clients and servers. A client and server are generally remote from each other and typically interact through a network, such as the described one. The relationship of client and server arises by virtue of computer programs running on the respective computers and having a client-server relationship to each other.

A number of implementations have been described. Nevertheless, it will be understood that various modifications may be made without departing from the spirit and scope of this disclosure. Accordingly, other implementations are within the scope of the following claims.

What is claimed is:

1. A computer-implemented method comprising:
  - receiving an input query that relates to multiple attribute classes of data records having attribute classes associated therewith in data storage;
  - identifying the multiple attribute classes related to the input query, the multiple attribute classes being a subset of less than all of the attribute classes associated with the data records in the data storage;
  - accessing, from electronic storage, an index that is descriptive of the data records in the data storage, the index defining positional relationships between the attributes classes of the data records;
  - identifying relative positions in the accessed index for the multiple attribute classes identified as being related to the input query;
  - modifying the accessed index based on the identified relative positions, wherein modifying the accessed index based on the identified relative positions comprises:
    - identifying portions of the accessed index that correspond to the multiple attribute classes identified as being related to the input query and removing portions of the accessed index that correspond to attribute classes other than the multiple attribute classes identified as being related to the input query;
    - based on the identified relative positions, selecting, from among the multiple attribute classes identified as being related to the input query, the attribute class that has the highest relative position within the accessed index; and
    - for each distinct value in the selected attribute class, assigning a distinct identifier to portions of the accessed index within the selected attribute class that correspond to the distinct value,

40

identifying, using the modified index, groups of data records that are included in the data records in the data storage and that share common values for the multiple attribute classes identified as being related to the input query;

accessing, from the data storage, metrics corresponding to the data records included in each of the identified groups of data records;

performing computations on the accessed metrics that distinguish among each of the identified groups of data records;

generating a report that reflects results of the computations; and

storing, in electronic storage, the generated report.

2. The method of claim 1 wherein modifying the accessed index based on the identified relative positions further comprises compressing redundancies in the accessed index that result from the identified relative positions.

3. The method of claim 1:

wherein modifying the accessed index based on the identified relative positions comprises:

identifying portions of the index that correspond to the multiple attribute classes identified as being related to the input query and disregarding portions of the accessed index that correspond to attribute classes other than the multiple attribute classes identified as being related to the input query; and

generating an updated index using the identified portions of the index that correspond to the multiple attribute classes identified as being related to the input query; and wherein identifying, using the modified index, groups of data records that are included in the data records in the data storage and that share common values for the multiple attribute classes identified as being related to the input query comprises identifying, using the updated index, groups of data records that are included in the data records in the data storage and that share common values for the multiple attribute classes identified as being related to the input query.

4. The method of claim 1 wherein modifying the accessed index based on the identified relative positions comprises:

building a connectivity graph based on the multiple attribute classes identified as being related to the input query and the identified relative positions; and

assigning group numbers to each node within the connectivity graph, each group number corresponding to a distinct combination of values for the multiple attribute classes identified as being related to the input query.

5. The method of claim 4 wherein identifying, using the modified index, groups of data records that are included in the data records in the data storage and that share common values for the multiple attribute classes identified as being related to the input query comprises identifying, using the connectivity graph and the assigned group numbers, groups of data records that are included in the data records in the data storage and that share common values for the multiple attribute classes identified as being related to the input query.

6. The method of claim 1 wherein accessing, from the data storage, metrics corresponding to the data records included in each of the identified groups of data records comprises:

automatically, without user intervention, generating database code for accessing data corresponding to the identified groups of data records; and

accessing the metrics using the generated database code.

7. The method of claim 1:

wherein identifying, using the modified index, groups of data records that are included in the data records in the

41

data storage and that share common values for the multiple attribute classes identified as being related to the input query comprises identifying addresses of data records mapped to by the processed index; and wherein accessing, from the data storage, metrics corresponding to the data records included in each of the identified groups of data records comprises accessing, from the data storage, metrics corresponding to the identified addresses of data records mapped to by the processed index.

8. The method of claim 1:

wherein performing computations on the accessed metrics that distinguish among each of the identified groups of data records comprises performing a computation for each group of the accessed metrics; and

wherein generating the report that reflects results of the computations comprises generating a report that includes the result of the computation for each group of the accessed metrics along with values for the multiple attribute classes identified as being related to the input query that correspond to the group of the accessed metrics.

9. The method of claim 1 wherein receiving the input query that relates to multiple attribute classes of data records having attribute classes associated therewith in data storage comprises receiving a report generation request that identifies multiple attribute classes to use in organizing the report.

10. An electronic system comprising:

at least one electronic data storage device; and

at least one processor configured to perform operations comprising:

receiving an input query that relates to multiple attribute classes of data records having attribute classes associated therewith in data storage;

identifying the multiple attribute classes related to the input query, the multiple attribute classes being a subset of less than all of the attribute classes associated with the data records in the data storage;

accessing, from electronic storage, an index that is descriptive of the data records in the data storage, the index defining positional relationships between the attributes classes of the data records;

identifying relative positions in the accessed index for the multiple attribute classes identified as being related to the input query;

modifying the accessed index based on the identified relative positions, wherein modifying the accessed index based on the identified relative positions comprises;

identifying portions of the accessed index that correspond to the multiple attribute classes identified as being related to the input query and removing portions of the accessed index that correspond to attribute classes other than the multiple attribute classes identified as being related to the input query;

based on the identified relative positions, selecting, from among the multiple attribute classes identified as being related to the input query, the attribute class that has the highest relative position within the accessed index; and

for each distinct value in the selected attribute class, assigning a distinct identifier to portions of the accessed index within the selected attribute class that correspond to the distinct value,

identifying, using the modified index, groups of data records that are included in the data records in the data

42

storage and that share common values for the multiple attribute classes identified as being related to the input query;

accessing, from the data storage, metrics corresponding to the data records included in each of the identified groups of data records;

performing computations on the accessed metrics that distinguish among each of the identified groups of data records;

generating a report that reflects results of the computations; and

storing, in electronic storage, the generated report.

11. The system of claim 10 wherein modifying the accessed index based on the identified relative positions further comprises compressing redundancies in the accessed index that result from the identified relative positions.

12. The system of claim 10:

wherein modifying the accessed index based on the identified relative positions comprises:

identifying portions of the index that correspond to the multiple attribute classes identified as being related to the input query and disregarding portions of the accessed index that correspond to attribute classes other than the multiple attribute classes identified as being related to the input query; and

generating an updated index using the identified portions of the index that correspond to the multiple attribute classes identified as being related to the input query; and wherein identifying, using the modified index, groups of data records that are included in the data records in the data storage and that share common values for the multiple attribute classes identified as being related to the input query comprises identifying, using the updated index, groups of data records that are included in the data records in the data storage and that share common values for the multiple attribute classes identified as being related to the input query.

13. The system of claim 10 wherein modifying the accessed index based on the identified relative positions comprises:

building a connectivity graph based on the multiple attribute classes identified as being related to the input query and the identified relative positions; and

assigning group numbers to each node within the connectivity graph, each group number corresponding to a distinct combination of values for the multiple attribute classes identified as being related to the input query.

14. The system of claim 13 wherein identifying, using the modified index, groups of data records that are included in the data records in the data storage and that share common values for the multiple attribute classes identified as being related to the input query comprises identifying, using the connectivity graph and the assigned group numbers, groups of data records that are included in the data records in the data storage and that share common values for the multiple attribute classes identified as being related to the input query.

15. The system of claim 10 wherein accessing, from the data storage, metrics corresponding to the data records included in each of the identified groups of data records comprises:

automatically, without user intervention, generating database code for accessing data corresponding to the identified groups of data records; and accessing the metrics using the generated database code.

16. The system of claim 10:

wherein identifying, using the modified index, groups of data records that are included in the data records in the data

data storage and that share common values for the multiple attribute classes identified as being related to the input query comprises identifying addresses of data records mapped to by the processed index; and  
wherein accessing, from the data storage, metrics corresponding to the data records included in each of the identified groups of data records comprises accessing, from the data storage, metrics corresponding to the identified addresses of data records mapped to by the processed index.

17. The system of claim 10:

wherein performing computations on the accessed metrics that distinguish among each of the identified groups of data records comprises performing a computation for each group of the accessed metrics; and  
wherein generating the report that reflects results of the computations comprises generating a report that includes the result of the computation for each group of the accessed metrics along with values for the multiple attribute classes identified as being related to the input query that correspond to the group of the accessed metrics.

18. The system of claim 10 wherein receiving the input query that relates to multiple attribute classes of data records having attribute classes associated therewith in data storage comprises receiving a report generation request that identifies multiple attribute classes to use in organizing the report.

\* \* \* \* \*